Optimal Dynamic Contract of Influence

By Yi Chen*

Draft: December 7, 2018

I study the optimal dynamic contract in a long-term principal-agent relationship, where the agent privately observes an evolving state but his preferences are state-independent. The principal commits to action flows based solely on the agent's reports. I show that communication is generically effective despite the misaligned preferences. Moreover, the optimal contract can stipulate actions that move in the opposite direction of the principal's ideal actions; a necessary and sufficient condition is provided. The principal is worse off over time in expectation, but the agent is not necessarily immiserated. The results apply to dynamic allocation problems such as capital budgeting.

JEL: D82, D83, D86

Keywords: Communication; Dynamic contract; Contrarian; Prudence; Diffusion

In organizations, the expert of relevant information is often not the decision-maker, and the effectiveness of information transmission is limited by conflict of interests. The self-interested informed agent could have a motive to misguide the uninformed principal, who takes actions based on the communication. In this paper, I investigate whether or not effective information transmission can be induced by a dynamic contract, and how to best elicit and utilize the private information.

As an example, the headquarters of a firm must decide how to allocate resources over time to a division manager, who is endowed with a unique product. From the

* Chen: Cornell University, 106 East Ave, Ithaca, NY, 14853, email: yi.chen@cornell.edu. I am grateful to Dirk Bergemann, Yeon-Koo Che, Chen Cheng, Eduardo Faingold, Simone Galperti, Yuan Gao, Marina Halac, Johannes Hörner, Ryota Iijima, Justin Johnson, Navin Kartik, Nicolas Lambert, Elliot Lipnowski, Chiara Margaria, Dmitry Orlov, Gregory Pavlov, Larry Samuelson, Ennio Stacchetti, Philipp Strack, Juuso Välimäki, Zhe Wang, Yiqing Xing, and John Zhu for extended discussions. I also thank the conference or seminar participants at Penn State University, Cornell University, Johns Hopkins University, University of Southern California, University of Arizona, International Conference on Game Theory at Stony Brook, Southern Economic Association Annual Conference, and Western University for helpful comments and suggestions. All errors are mine.

perspective of the headquarters, the optimal amount of resources to be allocated to the division depends on some state of the product, say, profitability or a technical parameter. However, the division manager knows the state much better due to specialization. Communication frictions arise if the manager always prefers more resources to less, *regardless* of the product's actual state. The headquarters commits to a dynamic rule of resource allocation based on the manager's reports. With these severely misaligned preferences, does the headquarters benefit at all from consulting the manager? If yes, does the headquarters necessarily allocate more resources when the state of the product is optimistic, and vice versa?

To answer these questions, I use the framework of dynamic principal-agent model. The agent privately observes a state process which evolves according to a Brownian motion with drift. The agent continuously reports to the principal, with the ability to manipulate the report by inflating or shading the true process at any time. The principal observes nothing but the agent's reports, and commits to a dynamic contract specifying actions over time based on the report history. The principal values information about the state: she suffers a quadratic cost in the gap between the action and her "target". The target is a function of the state, specifying the ideal action of the principal given the state. In contrast, the agent is severely biased in that he always prefers higher action levels regardless of the state.

State-independent preferences of the agent can arise in a number of situations. An investment manager with empire-building motives always prefers more resources; a technician with emotional attachment to his own lifetime project invariably craves for more budget. It is well-known that communication is extremely difficult if the agent has state-independent preferences and the principal-agent relationship lasts for only one period. Indeed, the only way for the agent to tell truth is to assign the same expected action for all reported states.

The prospect for communication is much brighter when it comes to long-term relationships with an evolving state. In order to induce truthful reports from the agent, the principal now has additional degrees of freedom to adjust the sequence of actions while still keeping the agent's continuation payoff independent of the current report. These degrees of freedom allows the principal to reallocate actions inter-temporally in her favor. This is how a dynamic contract improves communication.

As a first result, I show that the optimal dynamic contract generically induces non-babbling communication, only except when the weighted discounted future target is always equal to the current target. In other words, the principal benefits from the inter-temporal trade-off while respecting the agent's incentives, as long as there *is* a profitable trade-off to make.

The intuition lies in how the target depends on the state, i.e. the shape of the target function. Dynamic optimality requires the principal to balance the present and future distortions. If the target is linear, then the principal finds herself in the spot where the current mismatch between the action and the target exactly equals the expected future mismatch, such that there is no direction for profitable trade-off and babbling is the best contract. Somewhat surprisingly, communication could fail even if the target is non-linear with some particular form, for the same reason.

More importantly, the second result shows that when communication is effective, the optimal contract could behave counter-intuitively: to increase the action when the target falls, and vice versa. This poses a stark contrast to the literature of allocation problems where a quota-like mechanism prescribes a high allocation at high states, at the cost of reducing the future pool of allocations. In this paper, it could be optimal to *save* the quota when the state demands high allocations, and *use* it otherwise. I find a simple necessary and sufficient condition for this to occur, which turns out to involve the third derivative of the target function.

Why is it ever optimal to act against the target, and why should the third derivative determine the nature of the contract? For simplicity, suppose the state process has no drift, and the target function is increasing with a positive third derivative. If there is an increase in the state, then the current target is higher, and naturally the principal is inclined to raise the action level. Meanwhile, due to the persistence of the state, future targets are higher in expectation too, and therefore the principal is also tempted to allocate higher actions in the future to bridge the gaps. Unfortunately, because of the agent's incentive constraints, the principal cannot achieve both. To induce truth-telling, she has to either increase the current action *and* decrease the future actions, or the other way around. Which way is better? Again, the shape of the target matters. Since the target has a positive third derivative, its slope is convex in the state. Due to the random state evolution, the expected slope in the future is higher than the current slope by Jensen's inequality. In other words, the future targets are more sensitive to the current increase of the state, and hence in the inter-temporal trade-off the principal optimally sacrifices the current matching quality in exchange for better matches in the future.

Moreover, the model delivers predictions regarding the long-term well-being of the two parties. The principal faces an ever-increasing cost on average, meaning that the quality of match is deteriorating over time. This is not attributed to the cost-backloading argument; rather, the incentive constraints cause distortions to accumulate over time. The agent, on the other hand, is not necessarily immiserated. His payoff may drift up or down without bound depending on the shape of the target function.

I explore two extensions of the main model. The first is mean-reverting state process. It turns out that some persistence in the state is necessarily for the action to move against the target. If the state displays strong mean reversion, then the future targets are insensitive to the current state change, and there is little reason to sacrifice the current action for future. This explains why such contracts seldom arise in allocation problems where the state takes only two values: mean-reversion is automatically built in the two-state Markov chain.

The second extension is to allow for monetary transfers with limited liability. Depending on the tightness of this constraint, the optimal contract behaves accordingly. When the relationship evolves to the point where the limited liability is almost binding, then the predictions of the main model is qualitatively preserved. Otherwise, the principal uses money as a cheaper method than allocation to discipline the agent.

This paper is connected to the literature of communication in general. Since Vincent P. Crawford and Joel Sobel (1982) and Jerry R. Green and Nancy L. Stokey (2007), there is a large body of literature on cheap talk with a fixed one-dimensional state (Robert J. Aumann and Sergiu Hart (2003), Vijay Krishna and John Morgan (2001, 2004), Maria Goltsman, Johannes Hörner, Gregory Pavlov and Francesco Squintani (2009), etc.). In these papers the static nature of decision requires significant congruence of preferences in order for informative equilibria to exist. Commitment power of the principal brings my paper closer to the literature of delegation (Bengt Robert Holmström (1977), Ricardo Alonso and Niko Matouschek (2008), Manuel Amador and Kyle Bagwell (2013)), but dynamic delegation with full commitment is relatively understudied.

The paper also belongs to the literature of multi-dimensional influence game or allocation problem. Improved communication is found in static cheap talk games with higher dimensional state (Marco Battaglini (2002), Archishman Chakraborty and Rick Harbaugh (2010)) or multi-stage cheap talk (Mikhail Golosov, Vasiliki Skreta, Aleh Tsyvinski and Andrea Wilson (2014)). Asymptotic efficiency is obtained in repeated cheap talk relationships (Jérôme Renault, Eilon Solan and Nicolas Vieille (2013), Chiara Margaria and Alex Smolin (2017)). Static or dynamic models of multiple competing agents (Nemanja Antič and Kai Steverson (2016), Daniel Fershtman (2017)) find "strategic favortism" as an optimal way to discipline agents. Andrey Malenko (2016) and Raphael Boleslavsky and Tracy R Lewis (2016) are dynamic mechanisms of influence with either costly verification or noisy observation of the state. Among this literature two papers are closely related. Matthew O Jackson and Hugo F Sonnenschein (2007) consider multiple allocations together, and introduce a quota mechanism that utilizes the Law of Large Numbers to achieve asymptotic efficiency. Due to its dynamic nature, my model also features a "quota" (continuation payoff of the agent) that links decisions across periods, but the usage of quota can be completely against the principal's temptation in some cases. Also, my model gives the exact optimal mechanism for a fixed discount rate, instead of asymptotic efficiency. Yingni Guo and Johannes Hörner (2017) characterizes the optimal dynamic mechanism without transfer, which features a generalized quota. My paper shares the feature that the agent's private information is serially correlated, and the conflicts of interest is severe. However, in my model the agent's payoff is state-independent and the state is more persistent than a two-state Markov chain, and consequently the qualitative nature of the optimal contract is different.

Last but not least, the paper benefits from the insight of the literature on dynamic agency problems with transfer (Ana Fernandes and Christopher Phelan (2000), Marco Battaglini (2005), Yuliy Sannikov (2008), Noah Williams (2011), Marek Kapička (2013), Peter M DeMarzo and Yuliy Sannikov (2016), Zhiguo He, Bin Wei, Jianfeng Yu and Feng Gao (2017)). The absence of transfers in my model generates quite different implications for the optimal contract.

The remainder of the paper is organized as follows. Section I presents a two-period example to illustrate the key trade-off in the optimal contract. Section II lays out the setting for the continuous-time model. Section III simplifies the problem through revelation principle and derives a necessary condition for incentive compatibility. Section IV fully analyzes the optimal contract and gives some applications. Section V discusses two extensions, and Section VI concludes.

I. A Two-Period Example

To illustrate inter-temporal trade-offs that shape a dynamic contract, it is convenient to start with a static example and then contrast it with a two-period example.

First suppose there is only one period. A random state θ has normal distribution $\mathcal{N}(0,1)$. The agent observes θ and reports $\hat{\theta} \in \mathbb{R}$ to the principal, who then takes an action $x \in \mathbb{R}$ according to a pre-committed mechanism $x = x(\hat{\theta})$. There is no transfer. Given the true state θ and the actual action x, the principal's cost is $(x - f(\theta))^2$, and the agent's payoff is simply x. In other words, while the agent prefers higher actions regardless of the state, the principal tries to match the action with some state-contingent target, $f(\theta)$. In this static environment with severe conflict of interests, meaningful communication cannot be induced in any mechanism. Indeed, a mechanism must map all states into the same expected action to respect the incentives of the agent. Because of the convex loss function, the principal cannot do better than committing to a constant action at $\mathbb{E}\theta = 0$, which is unresponsive to the reported state.

The hope of meaningful communication is not entirely lost, however. Let us move a step forward to the simplest dynamic example, with two periods $t = \{1, 2\}$. The state θ evolves over time according to random walk:

$$\theta_1 = \varepsilon_1, \quad \theta_2 = \theta_1 + \varepsilon_2,$$

where ε_1 and ε_2 are independently drawn from the normal distribution $\mathcal{N}(0,1)$. In each period t, the agent observes θ_t and reports $\hat{\theta}_t \in \mathbb{R}$ to the principal, who then relies solely on the report history to take action x_t and ends period t. Again, there is no transfer. The principal's total cost from both periods is $(x_1 - f(\theta_1))^2 + (x_2 - f(\theta_2))^2$, and the agent's total payoff is $x_1 + x_2$. Notice that the target function $f(\cdot)$ is invariant to time. A dynamic contract is a pair $(x_1(\hat{\theta}_1), x_2(\hat{\theta}_1, \hat{\theta}_2))$, mapping report histories into actions. Without loss of generality I focus on contracts that induce on-path truth-telling for the agent.

To obtain the optimal mechanism, the principal solves the following:

$$\min_{x_1(\cdot),x_2(\cdot,\cdot)} \quad \mathbb{E}\left[(x_1 - f(\theta_1))^2 + (x_2 - f(\theta_2))^2\right]$$

$$\text{(1)} \quad \text{s.t.} \quad x_1(\theta_1) + \mathbb{E}\left[x_2(\theta_1,\theta_2)|\theta_1\right] \geqslant x_1(\hat{\theta}_1) + \mathbb{E}\left[x_2(\hat{\theta}_1,\hat{\theta}_2)\middle|\theta_1\right] \quad \forall\ \theta_1,\hat{\theta}_1,\hat{\theta}_2,$$

$$\text{(2)} \quad x_2(\theta_1,\theta_2) \geqslant x_2(\theta_1,\hat{\theta}_2) \quad \forall\ \theta_1,\theta_2,\hat{\theta}_2.$$

Constraint (2) requires that truth-telling is optimal for the agent in period 2 after a truthful report in period 1. Constraint (1) governs the period-1 incentive, stating that the agent obtains the highest expected total payoff by truth-telling in both periods, among all reporting strategies.

Since the agent's payoff is completely state-independent, condition (2) implies $x_2(\theta_1,\theta_2) = x_2(\theta_1)$ for all θ_2. That is, x_2 must be independent of θ_2, and I abuse notation by writing $x_2(\theta_1)$ for short. Moving back to period 1, condition (1) is then simplified to $x_1(\theta_1)+x_2(\theta_1) \geqslant x_1(\hat{\theta}_1)+x_2(\hat{\theta}_1)$ for all θ_1 and $\hat{\theta}_1$. By switching the pair of states, condition (1) reduces to:

$$x_1(\theta_1) + x_2(\theta_1) \equiv W,$$

where W on the right-hand side is a constant, naturally interpreted as the (fixed) total payoff of the agent. The optimal level of W is endogenously chosen by the principal as part of the maximization problem.

At this point, the role of dynamics becomes clear. In the one-period case, incentive compatibility requires the contract to specify an action independent of the reported state, and therefore information is wasted and communication fails. The logic is different when there are two periods. Indeed, period-2 report is again ignored, reducing communication to babbling in that period. However, the period-1 incentive constraint only requires the *total payoff* of the agent to be unresponsive to $\hat{\theta}_1$, but the *action pair* (x_1, x_2) still has one degree of freedom to adjust. While the agent is indifferent among all action pairs that have the same sum, the principal, who has different preferences, values the ability to reallocate actions between periods in response to period-1 information.

The optimal contract reads (derivation relegated to Appendix):

$$x_1^*(\theta_1) = \frac{1}{2}\left(f(\theta_1) - \mathbb{E}[f(\theta_2)|\theta_1]\right) + \frac{1}{2}\left(\mathbb{E}f(\theta_1) + \mathbb{E}f(\theta_2)\right), \tag{3}$$

$$x_2^*(\theta_1) = \underbrace{-\frac{1}{2}\left(f(\theta_1) - \mathbb{E}[f(\theta_2)|\theta_1]\right)}_{\text{responsive to } \theta_1} + \underbrace{\frac{1}{2}\left(\mathbb{E}f(\theta_1) + \mathbb{E}f(\theta_2)\right)}_{\text{independent of } \theta_1}. \tag{4}$$

To interpret, both x_1^* and x_2^* can be decomposed into a term responsive to θ_1, and a constant term. The pair of actions add up to $W^* \equiv x_1^*(\theta_1) + x_2^*(\theta_1) = \mathbb{E}f(\theta_1) + \mathbb{E}f(\theta_2)$, meaning that the optimal choice of the total payoff is the unconditional expectation of total targets across periods.[1] Also, for either period, the constant term is simply half of W^*, a result from "cost smoothing".

More importantly, the first terms in x_1^* and x_2^* respond to the period-1 report. The period-1 incentive constraint $x_1 + x_2 = W$ serves as a "budget" or "quota" for inter-temporal allocation of actions. Along this budget line, x_1^* and x_2^* reacts to period-1 report for cost minimization. To be precise, differentiating both sides of (3) with respect to θ_1 (assuming existence of f'), we have:

$$\frac{\mathrm{d}x_1^*}{\mathrm{d}\theta_1} = \frac{1}{2}\left(f'(\theta_1) - \frac{\mathrm{d}}{\mathrm{d}\theta_1}\mathbb{E}\left[f(\theta_2)|\theta_1\right]\right) = \frac{1}{2}\left(f'(\theta_1) - \mathbb{E}\left[f'(\theta_2)|\theta_1\right]\right), \tag{5}$$

where the second equality results from the random walk process. The sign of this derivative is ambiguous. On one hand, a marginal change in the state θ_1 calls for a corresponding change in the action x_1 to better match the state, hence the term $f'(\theta_1)$. On the other hand, the marginal change in θ_1 also forecasts an expected change in θ_2, demanding an adjustment of action in period 2. Due to the inter-temporal "budget constraint," any adjustment of x_2 is at the cost of x_1, causing an offsetting term $-\mathbb{E}\left[f'(\theta_2)|\theta_1\right]$ in the bracket in (5). The relative sizes of these competing effects shape the dynamic contract.

If the first effect is stronger, then it is more cost-efficient to let the action x_1^* move in the *same* direction as the target $f(\theta_1)$, at the cost of worse expected match in period 2. The action is thus called *conformist* at state θ_1. If the second effect is stronger, then cost-efficiency requires the action x_1^* to move in the *opposite* direction as the target as a sacrifice, in exchange for a better match in period 2. If this is the case, then the action is called *conformist* at state θ_1. For

[1] This nice feature relies on the quadratic cost structure.

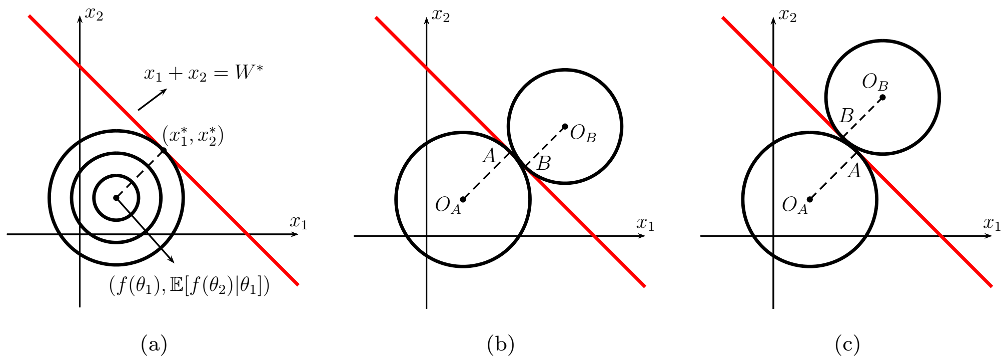

FIGURE 1. ISO-COST CURVES AND THE BUDGET LINE. PANEL (A): THE FAMILY OF ISO-COST CURVES AND THE BUDGET LINE. PANEL (B): CONFORMIST—x_1 AND $f(\theta_1)$ MOVE IN THE SAME DIRECTION. PANEL (C): CONTRARIAN—x_1 AND $f(\theta_1)$ MOVE IN OPPOSITE DIRECTIONS.

instance, when $f(\theta) = e^{\theta}$, the action is contrarian at all θ_1. For target function $f(\theta) = \theta - \frac{1}{2}|\theta|$, contrarian action occurs only at a subset of states. The action is always conformist if $f(\theta) = \theta\sqrt{|\theta|}$. Moreover, if $f(\theta) = \theta$ or $f(\theta) = \theta^2$, then the two terms in the bracket of (5) always cancel out. As a result, the action pair does not respond to information at all, indicating a failure of communication.

Figure 1 geometrically illustrates the cost minimization problem. In a (x_1, x_2)-plane, incentive compatibility pins the action pair on a budget line with slope -1, shown as the red line in Panel (a). The concentric circles are the principal's iso-cost curves. The center of the circles has coordinates $(f(\theta_1), \mathbb{E}[f(\theta_2)|\theta_1])$, namely, the period-1 target and the conditional expected period-2 target. Both are determined by θ_1 only. The constrained optimal choice of actions is found on the tangent point (x_1^*, x_2^*), which spans a 45-degree line from the center. Panel (b) shows conformist actions. Suppose a change in θ_1 causes the current target $f(\theta_1)$ to increase, it also leads to a higher expected target $\mathbb{E}[f(\theta_2)|\theta_1]$ in period 2, due to persistence of the state. Graphically, the change in θ_1 shifts the center of circles in the northeast direction from O_A to O_B, but the horizontal shift dominates (O_B lies below the 45-degree line $O_A A$). In this case, the new contract (tangent point B) asks for a higher current action x_1 than before (point A). Panel (c) depicts another possibility: the increase in expected period-2 target is more significant than the increase in the current target. The new contract allocates *lower* current action x_1, against the agent's report, in order to save for a higher action in period 2 for a better matching.

In the next section, I lay out the formal model in continuous time and with infinite horizon. Continuous time allows for the smooth evolution of information and closed-form analysis. Infinite horizon enables the study of the asymptotics of the contract.

II. The Continuous-Time Model

There is a principal (she) and an agent (he). Time $t \geqslant 0$ is continuous. A stochastic process $\theta = (\theta_t)_{t\geqslant 0}$, called the *state*, evolves according to:

$$\theta_t = \theta_0 + \mu t + \sigma Z_t,$$

where $Z = (Z_t)_{t\geqslant 0}$ is the standard Brownian motion on the probability space $(\Omega, \mathscr{F}, \mathbb{P})$. Constants μ and $\sigma > 0$ are drift and volatility, respectively. The initial state θ_0 is common knowledge.

Over time, the agent reports a *manipulated version* $\hat{\theta} = (\hat{\theta}_t)_{t\geqslant 0}$ of the state process. Specifically:

$$\mathrm{d}\hat{\theta}_t = m_t \mathrm{d}t + \mathrm{d}\theta_t,$$

where m_t, chosen by the agent at every moment $t \geqslant 0$, is interpreted as the "intensity of manipulation." In other words, the reported process can drift away from the true process. The principal takes action $x_t \in \mathbb{R}$ at all times.

Interests are severely misaligned. While the principal's favorite action depends on the state, the agent only wishes to induce actions as high as possible. Specifically, the principal's flow cost at time t from a state-action pair (θ_t, x_t) is $(x_t - f(\theta_t))^2$, i.e., she suffers a quadratic cost from the gap between the actual action x_t and the state-dependent *target* action $f(\theta_t)$.[2] The agent's flow payoff is simply x_t, independent of the state.[3] The linear payoff of the agent shuts down risk aversion as a possible channel of rent extraction, as has been discovered in the literature.

The players discount future at rate $r > 0$. Given realized paths of action and state, $(x_t, \theta_t)_{t\geqslant 0}$, the total cost of the principal and total payoff of the agent are,

[2]The quadratic assumption is for simplicity. It is not essential for the qualitative results.

[3]This insatiable preference of the agent can be interpreted as taking the bias $b \to \infty$ in Vincent P. Crawford and Joel Sobel (1982), although the main results do not rely on this extreme specification.

respectively:

$$u_P\left((x_t,\theta_t)_{t\geqslant 0}\right) = \int_0^\infty re^{-rt}(x_t - f(\theta_t))^2 \mathrm{d}t,$$
$$u_A\left((x_t,\theta_t)_{t\geqslant 0}\right) = \int_0^\infty re^{-rt}x_t \mathrm{d}t,$$

whenever well-defined.

Other than the initial state θ_0, the principal's only information about the state is the agent's report history. That is, the principal does not observe her own flow payoffs or any signals about past states.

The principal commits to a contract at time zero. I use superscript t to denote a history up to time t. A contract x is a $\hat{\theta}$-measurable process specifying an action $x_t(\hat{\theta}^t) \in \mathbb{R}$ as a function of the report history, for all $t \geqslant 0$. There are no monetary transfers.[4] A strategy of the agent is a θ-measurable process m. It prescribes the drift $m_t(\theta^t)$ that the agent adds to the true state process as a function of the state history, for all $t \geqslant 0$. I define the space of feasible strategies as:

$$\mathcal{M} \equiv \left\{ m : \mathbb{E}\left[e^{2\overline{\alpha}\int_0^t m_s \mathrm{d}s}\right] < \infty\ \forall t,\ \text{and}\ \lim_{t\to\infty} e^{-rt}\mathbb{E}\left[e^{2\overline{\alpha}\int_0^t m_s \mathrm{d}s}\right] = 0 \right\},$$

where $\overline{\alpha} \equiv \frac{\sqrt{2r\sigma^2+\mu^2}-|\mu|}{2\sigma^2}$, to exclude explosive strategies. This restriction is not essential; In the end of Appendix I show the effect of relaxing the strategy set.

Given a contract-strategy pair (x, m), the total expected cost and payoff are respectively:

$$U_P(x,m) = \mathbb{E}^m\left[\int_0^\infty re^{-rt}(x_t - f(\theta_t))^2 \mathrm{d}t\right],$$
$$U_A(x,m) = \mathbb{E}^m\left[\int_0^\infty re^{-rt}x_t \mathrm{d}t\right],$$

whenever well-defined, where $\mathbb{E}^m$ denotes the expectation induced by strategy m. Hereafter, "payoff" and "cost" refer to the agent's total expected payoff and the principal's total expected cost, unless otherwise noted.

The agent chooses a strategy m to maximize his payoff given contract x. The principal designs a contract x to minimize her cost given the agent's optimal

[4]This is often the case in many organizations. Even if there are monetary transfers, I show in Section V.B that as long as there is limited liability for the agent, the qualitative results are preserved when the limited liability is close to be binding.

choice of strategy in reaction to the contract.

III. Incentives of the Agent

This section performs two standard steps of simplification, facilitating the analysis of the optimal contract in Section IV. First, I restrict attention to truthful contracts by invoking a version of the Revelation principle. Second, I use the first-order approach to derive a necessary condition for incentive compatibility. Sufficiency is verified after the candidate solution is obtained in Section IV.

A. Revelation Principle

A strategy m is called *truthful* if it is identically zero, denoted as $m^\dagger$. A contract x is called *truthful* if the agent maximizes his payoff with the truthful strategy. By Lemma 1 below, I can focus on truthful contracts without loss of generality.

LEMMA 1 (Revelation Principle):
Given any contract x that implements a mapping from state paths into action paths, there exists a truthful contract $x^\dagger$ that implements the same mapping.

Among truthful contracts (hereafter "truthful" is by default), the principal solves:

$$\min_{(x_t(\cdot))_{t\geqslant 0}} \quad \mathbb{E}\left[\int_0^\infty re^{-rt}(x_t(\theta^t)-f(\theta_t))^2\mathrm{d}t\right] \tag{6}$$

$$\text{s.t.} \quad \mathbb{E}\left[\int_0^\infty re^{-rt}x_t(\theta^t)\mathrm{d}t\right] \geqslant \mathbb{E}\left[\int_0^\infty re^{-rt}x_t(\hat{\theta}^t)\mathrm{d}t\right],$$

$$\text{where } \hat{\theta}_t = \theta_t + \int_0^t m_s\mathrm{d}s,\ \forall\, m \in \mathcal{M}. \tag{7}$$

The incentive constraint (7) guarantees that any strategy achieves at best the payoff from truth-telling. While the constraint is expressed as of time zero, it also implies incentive compatibility at all later times since the agent faces a decision problem with time-consistent preferences. Hidden behind (7) is the premise that the payoff of the agent is well-defined. This is without loss of generality, because if a contract generates non-integrable payoffs for the agent, it must bring infinite cost to the principal, which is clearly suboptimal.

B. Incentive Compatibility: Necessary Condition

The dynamic version of the first-order approach (Noah Williams (2011), Marek Kapička (2013), Alessandro Pavan, Ilya Segal and Juuso Toikka (2014), Peter M DeMarzo and Yuliy Sannikov (2016)) derives a local version of the incentive constraints, namely, conditions to prevent the agent from local deviations.

To apply this method, I define a process $W = (W_t)_{t \geqslant 0}$ for any contract x, by:

$$W_t(x) \equiv \mathbb{E}_t \left[\int_t^\infty r e^{-r(s-t)} x_s \mathrm{d}s \right],$$

as the agent's on-path expected continuation payoff. The expectation $\mathbb{E}_t$ is conditional on the information generated by the state by time t. As is verified later, this W_t together with θ_t suffices to summarize the history up to t.[5] From now on, I suppress the dependence of W on x to simplify notations.

Given a contract x, the evolution of W can be written as a diffusion process according to Lemma 2.

LEMMA 2 (Martingale Representation Theorem):
For any contract x, there exists a $\hat{\theta}$-measurable process $\beta = (\beta_t)_{t \geqslant 0}$ such that:

$$\mathrm{d}W_t = r(W_t - x_t)\mathrm{d}t + r\beta_t \underbrace{(\mathrm{d}\hat{\theta}_t - \mu \mathrm{d}t)}_{=\sigma \mathrm{d}Z_t}. \tag{8}$$

The first term on the right-hand side of (8)represents the drift of W_t: it grows at rate r and depletes as x_t is paid out to the agent. The second term, which is the diffusion, governs the incentives. On equilibrium path, it holds that $\mathrm{d}\hat{\theta}_t - \mu \mathrm{d}t = \sigma \mathrm{d}Z_t$, which has zero mean. The multiplier $r\beta_t$ is interpreted as the instantaneous slope of the continuation payoff with respect to reported states, or "strength of incentives" (Zhiguo He, Bin Wei, Jianfeng Yu and Feng Gao (2017), Peter M DeMarzo and Yuliy Sannikov (2016)). Suppose the instantaneous slope $r\beta_t$ is positive. By adding a drift $m > 0$ to the report for a short moment $\mathrm{d}t$, the

[5] The use of the continuation payoff as one of the sufficient statistics for the history is common when it comes to equilibrium payoffs (Dilip Abreu, David Pearce and Ennio Stacchetti (1986), Jonathan Thomas and Tim Worrall (1990), etc.), but in a setting with *persistent* private information, an additional state variable is often needed (Noah Williams (2011), Marek Kapička (2013), Yingni Guo and Johannes Hörner (2017), etc.). In my model, the marginal continuation payoff does not appear as a state variable even with persistent private information, because the agent's payoff is independent of the state. Specifically, the flow payoff and the evolution of the continuation payoff depend only on the action, which is publicly observed. When the agent lies, the perception of the agent's continuation payoff from the two parties coincide, even if they hold different beliefs about the state.

agent tricks the principal into believing that the state is $m\mathrm{d}t$ higher than it actually is, and this results in an increase of $r\beta_t m\mathrm{d}t$ in the agent's continuation payoff.[6] Similarly, if $r\beta_t < 0$, the agent can profit by shading the report. The only way to deter a local deviation from truth-telling is to keep $r\beta_t$ identically at zero. Proposition 1 formalizes the above reasoning as a necessary condition for incentive compatibility.

PROPOSITION 1 (IC-FOC):
A necessary condition for incentive compatibility is $\beta_t = 0$ for all $t \geqslant 0$.

According to the proposition, the only way to induce truth-telling from the agent is to entirely disentangle his continuation payoff from the current report, a theme already foreshadowed in the two-period model. While it may seem severe as a constraint for the principal, there are lots of degrees of freedom to maneuver: a given continuation payoff can be supported by infinitely many paths of actions. Therefore, the choice of how the action paths responds to information, subject to the IC-FOC, is the key to optimal utilization of information, and that is the task of Section IV.

IV. Optimal Contract

In this section I find the optimal contract and derive its properties. When using IC-FOC in place of the full incentive constraints, the solution is called a "candidate" until verified to satisfy the original constraints.

From now on, I impose some regularity conditions on the target function.

ASSUMPTION 1 (Regularity):
(i) The target f is piecewise $\mathcal{C}^2$;
(ii) There exists $\alpha_0 > 0, \alpha_1 \in [0, \overline{\alpha})$ such that $|f(\theta)| \leqslant \alpha_0(e^{\alpha_1\theta} + e^{-\alpha_1\theta})$.

Part (i) of the assumption puts some smoothness on the target function to enable local analysis. Part (ii) uniformly bounds the target by some exponential function with a low growth rate. It ensures the target does not drift to infinity too fast, necessary for costs and payoffs to be finite. The definition of $\overline{\alpha}$ is given in Section II along with the strategy set $\mathcal{M}$.

[6] With continuous time, the importance of the current flow payoff is negligible.

A. Solving for the Optimal Contract

I use a recursive formulation to solve the relaxed problem where only the IC-FOC is considered. As is conjectured in Section III.B, the optimal contract can be written in terms of two state variables: the state θ and the continuation payoff W. In Theorem 1 I formally prove that the candidate contract derived in the recursive problem is indeed the solution to the original problem (6)-(7).

Define $C(\theta, W)$ as the cost function of the principal, which has the two arguments as conjectured. According to (8) and Proposition 1, the second argument evolves as $\mathrm{d}W_t = r(W_t - x_t)\mathrm{d}t$. Hence, the cost function must satisfy the following functional equation:

$$\begin{aligned} rC(\theta, W) &= \min_x \; r(x - f(\theta))^2 \\ &\quad + r(W - x)C_W(\theta, W) + \mu C_\theta(\theta, W) + \frac{\sigma^2}{2} C_{\theta\theta}(\theta, W). \end{aligned} \tag{9}$$

The right-hand side of the equation consists of four terms. The first term is the normalized flow cost. The second and third terms are expected changes in cost caused by the drift in W and θ. The last term, the Itô term, reflects the volatility in θ. In order for the recursive form to make sense, the cost and the payoff must also satisfy the transversality condition:

$$\lim_{t\to\infty} e^{-rt}\mathbb{E}\left[C(\theta_t, W_t)\right] = 0, \;\; \lim_{t\to\infty} e^{-rt}\mathbb{E}W_t = 0. \tag{10}$$

Conditions (9) and (10) admit a candidate solution for the cost and policy functions:

$$C(\theta, W) \equiv (W - \gamma \star f(\theta))^2 + \frac{\sigma^2}{r}\gamma \star (\gamma \star f)'^2(\theta), \tag{11}$$

$$x(\theta, W) \equiv f(\theta) - \gamma \star f(\theta) + W, \tag{12}$$

where γ is a kernel with asymmetric Laplace distribution:

$$\gamma(z) \equiv \frac{r}{\sqrt{\mu^2 + 2r\sigma^2}} e^{\frac{\mu}{\sigma^2}z - \frac{\sqrt{\mu^2+2r\sigma^2}}{\sigma^2}|z|},$$

and $\gamma \star f$ is the convolution between the two functions. Panel (a) of Figure 2 shows the graph of the kernel γ for different values of the drift μ. When $\mu = 0$,

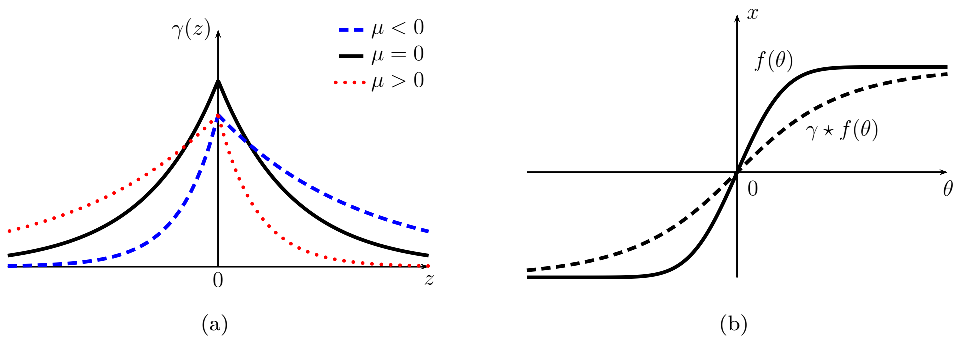

FIGURE 2. THE KERNEL AND THE CONVOLUTION. PANEL (A): GRAPHS OF THE KERNEL γ FOR $\mu < 0$, $\mu = 0$, AND $\mu > 0$. PANEL (B): THE TARGET FUNCTION f (SOLID CURVE) AND THE CONVOLUTION $\gamma \star f$ (DASHED CURVE) FOR $\mu = 0$.

the Laplace distribution is symmetric, otherwise it is skewed according to the sign of the drift. The convolution $\gamma \star f$ is economically interpreted as the *expected discounted future target*, which summarizes global information about the target function with weights depending on the state process. Indeed, it can be shown by Fubini Theorem that $\gamma \star f(\theta) = \mathbb{E}\left[\int_0^\infty re^{-rt} f(\theta_t)\mathrm{d}t \middle| \theta_0 = \theta\right]$. Panel (b) of Figure 2 presents a typical target function and its convolution with γ, when $\mu = 0$.

Theorem 1 below verifies that (11) and (12) achieve the minimum cost in the original problem (6)-(7).

THEOREM 1 (Optimal Contract):
The principal's minimum cost is $\frac{\sigma^2}{r}\gamma \star (\gamma \star f)'^2(\theta_0)$, *achievable with the essentially unique optimal contract:*

$$(13) \quad x_t(\theta^t) \quad \equiv \quad f(\theta_t) - \gamma \star f(\theta_t) + \underbrace{\gamma \star f(\theta_0)}_{W_0} + \underbrace{r \int_0^t (\gamma \star f(\theta_s) - f(\theta_s))\,\mathrm{d}s}_{W_t - W_0}.$$

Theorem 1 has several implications. First, since the principal is free to choose any initial W_0, the cost is minimized at $W_0 = \gamma \star f(\theta_0)$, which is the expected discounted future targets. That is, the principal should not distort the actions *on average* in the optimal contract; instead, the optimality comes from the way in which the action responds to the state. This will be discussed in detail in Subsection IV.B.

Second, the minimized cost is non-negative, and this cost is entirely due to agency problem. Indeed, without private information or preference misalignment, the cost could have been zero.

Third, the optimal contract echoes that of the two-period example: the action cashes out the agent's continuation payoff smoothly, and responds to the current state in two competing ways. Rearranging terms in the policy function, one arrives at the Euler's equation:

$$\underbrace{x(\theta, W) - f(\theta)}_{\text{current distortion}} \quad = \quad \underbrace{W - \gamma \star f(\theta)}_{\text{future distortion}} \ ,$$

which is more intuitive. The left-hand side is the gap between the current action x and the current target $f(\theta)$, or simply, the *current* distortion. The right-hand side is the gap between the expected discounted future action W and the expected discounted future target $\gamma \star f(\theta)$, or, the *future* distortion. The distortions must be balanced across time at optimum, meaning that it is optimal to set the current distortion as a fixed share of the future distortion. This smoothing motive is built in the convex flow cost of the principal.

The optimal contract sheds light on the effectiveness of communication from the principal's perspective. On one hand, the cost does not surpass the *babbling cost*, i.e., the lowest cost if the principal uses a deterministic action path. On the other, the cost is bounded below by zero—the cost as if there is complete information. In fact, the shape of the target function determines when the two bounds are achieved, as is summarized in Theorem 2.

THEOREM 2 (Impossibility):
(i) Zero cost is obtained if and only if the target is almost everywhere identical to a constant.
(ii) For $\mu = 0$, the cost reaches the babbling cost if and only if the target is almost everywhere identical to a polynomial of order up to 2.
(ii) For $\mu \neq 0$, the cost reaches the babbling cost if and only if the target is almost everywhere identical to $c_0 + c_1\theta + c_2 e^{-\frac{2\mu}{\sigma^2}\theta}$ for some constants c_0, c_1, c_2.

The theorem is interpreted as two impossibility results. First, friction-less communication is not achievable unless the state has no bearing on the principal's cost. Second, for some non-generic set of target functions, communication fails even with a dynamic contract. Specifically, if the target is affine in the state, then the action is a constant over time and babbling is the inevitable outcome.

Stranger still, babbling is optimal even when the target has some particular form of curvature. Theorem 2 conveys the message that it is necessary to investigate higher order derivatives of the target than simply the curvature. This is explored further in the next subsection.

B. *Conformist or Contrarian Response?*

With the optimal contract at hand, it is time to answer the questions raised in the Introduction: Should the action always move in the same direction as the target? When is it optimal to act "against" the agent's report? The following analysis provides necessary and sufficient conditions for the action to be conformist or contrarian. To proceed, I first formalize these two terms.

DEFINITION 1 (Contrarian vs Conformist):
At state θ where $f'(\theta)$ exists, the action x is called conformist *(*contrarian*, resp.) if:*

$$f'(\theta)\frac{\partial x}{\partial \theta} > 0 \ (< 0, \ \textit{resp.}),$$

i.e., the action moves in the same (opposite resp.) direction as the target.

The above definition makes sense because, as is explained below, in the optimal contract the conformist or contrarian property depends only the current state, not the state history.

Theorem 3 below proposes two interchangeable criteria to check in which direction the optimal action should move. The first criterion (ii) involves a comparison between f' and $(\gamma \star f)'$. The second criterion (iii) simply rewrites the same expression in terms of higher order derivatives.

THEOREM 3 (Conditions for Conformist and Contrarian):
The following statements are equivalent:
(i) The optimal contract stipulates a conformist (contrarian, resp.) action at state θ.
(ii) $(f'(\theta) - (\gamma \star f)'(\theta))f'(\theta) > 0 \ (< 0, \ \textit{resp.})$.
(iii) $\Big((\gamma \star f)''(\theta), (\gamma \star f)'''(\theta)\Big) \cdot (2\mu, \sigma^2) f'(\theta) < 0 \ (> 0, \ \textit{resp.})$.

According to the theorem, contrarian actions can occur in many scenarios. Subsection IV.D is dedicated to enumerating economically meaningful examples where contrarian actions arise. Why would contrarian action ever be optimal?

After all, it is tempting to increase the action following a rise in the target, the usual way that the "quota" of actions is used in the literature.

Here is the intuition. Heuristically, the slope $f'(\theta)$ can be interpreted as the *information sensitivity* of the principal at state θ. A steeper slope corresponds to a target that is more sensitive to a state change. Suppose $f' > 0$ for simplicity. If the current state has a shock $\mathrm{d}\theta > 0$, then the current target increases by $f'(\theta)\mathrm{d}\theta$, and the principal is tempted to increase the current action. Meanwhile, because of the persistence of the state process, the expected discounted future target increases as well, by $(\gamma \star f)'(\theta)\mathrm{d}\theta$. This also creates the motive to increase total future actions to better match the higher expected targets. However, she cannot achieve both due to incentive constraints: higher current action must be followed with lower future actions, and vice versa. Therefore, if $f'(\theta)$ is larger, then the current information sensitivity is higher than future, and the intertemporal trade-off tips the scale in favor of the current action. In this way, the current action increases along with the current target (conformist), at the cost of future matching. Conversely, if the future information sensitivity $(\gamma \star f)'(\theta)$ is higher, then the principal optimally sacrifices the current matching in order to rebalance the Euler equation. As a result, the current action moves against the current target (contrarian).

Either way, the ability to trade-off between the present and future benefits the principal, and this is how a dynamic contract facilitates communication above the babbling level. In the knife-edge case where $f' = (\gamma \star f)'$ for all states, there is no direction for profitable trade-off, and the principal is stuck with babbling. As is evident from Theorem 2, this can happen even if f is non-linear. For example, suppose $\mu = 0$ and f is quadratic. The information sensitivity f' is not a constant but is affine. Hence, when computing the expected future information sensitivity, the diffusion of the state cancel out and one arrives at the current information sensitivity. Only when f' itself has curvature, can the state uncertainty generate trade-off opportunities.

The optimal contract works in the similar pattern as the "quota mechanism" in the literature of allocation problems (Matthew O Jackson and Hugo F Sonnenschein (2007), Jérôme Renault, Eilon Solan and Nicolas Vieille (2013), Yingni Guo and Johannes Hörner (2017)). However, the important difference lies in *how* the quota is used. The usage of quota is "conformist" in the literature: as long as the quota is not depleted, spend it when the state is worth spending, and save it otherwise. In my paper, however, the optimal action can be contrarian depending on two factors: the shape of the target function and the (persistent) state process.

When is contrarian action more likely to occur? Condition (iii) of Theorem 3 gives a neat breakdown. Ignoring $f'(\theta)$, the left-hand side appears as the inner product of two vectors. The former characterizes the expected future target in terms of *absolute prudence* $-\frac{(\gamma\star f)'''(\theta)}{(\gamma\star f)''(\theta)}$, and the latter summarizes the state process through *normalized drift* $\frac{2\mu}{\sigma^2}$. This inner product boils down to two additive terms:

$$\underbrace{2\mu(\gamma\star f)''(\theta)}_{\text{drift effect}}+\underbrace{\sigma^2(\gamma\star f)'''(\theta)}_{\text{volatility effect}}.$$

If the expected future target is convex, i.e., $(\gamma\star f)''>0$, then with a positive drift the information sensitivity is increasing with time in expectation. This contributes to a more contrarian action through the drift effect when $f'>0$. On the other hand, if the expected future target has positive third derivative ($(\gamma\star f)'''>0$), then due to the diffusion ($\sigma^2>0$) the information sensitivity tends to increase over time by Jansen's inequality. This favors a contrarian action through the volatility effect when $f'>0$. The intuition for the case $f'<0$ is similar.

C. Evolution of the Contract

Over time, the cost and payoff evolve stochastically on the equilibrium path. In Proposition 2, I briefly explore the evolution of the contract to see the trends of the principal's costs and the agent's continuation payoff as the relationship moves on.

PROPOSITION 2 (Cost and Payoff Dynamics):
(i) The continuation cost is a sub-martingale, i.e., $\frac{\mathbb{E}_t\left[\mathrm{d}C_t\right]}{\mathrm{d}t}\geqslant 0$*;*
(ii) The continuation payoff monotonically increases (decreases, resp.) over time if $\gamma\star f(\theta)-f(\theta)>0$ *(*<0*, resp.), or equivalently* $2\mu\gamma\star f(\theta)+\sigma^2(\gamma\star f)'(\theta)>0$ *(*<0*, resp.), for all* θ*.*

Part (i) of Proposition 2 claims that the principal faces higher and higher costs in expectation as the contract is carried out over time. While the increasing cost is usually attributed to the distortion back-loading motive of the principal, it is not the case here. The actual reason is that with incentive constraints, the continuation payoff that the principal commits to the agent diverges from the

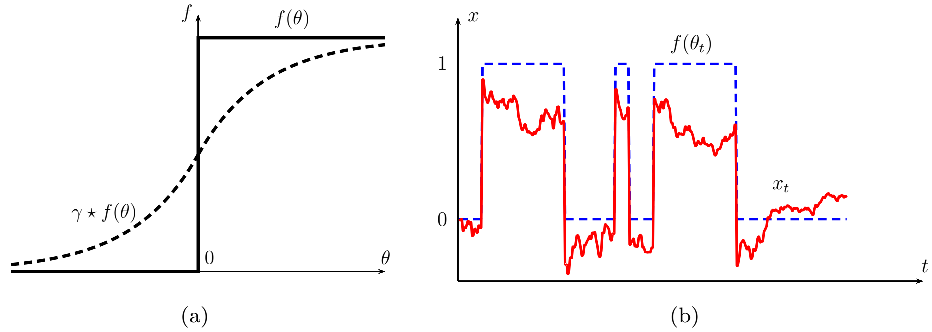

FIGURE 3. BINARY TARGET WITH PARAMETERS: $r = 1, \sigma = 1, \mu = 0, \theta_0 = -0.4$. PANEL (A): THE TARGET FUNCTION (SOLID CURVE) AND ITS CONVOLUTED VERSION (DASHED CURVE). PANEL (B): SIMULATED PATHS OF THE TARGET (BLUE DASHED CURVE) AND THE OPTIMAL ACTION (RED SOLID CURVE). THE ACTION ALWAYS JUMPS ONE-FOR-ONE WHEN THE TARGET JUMPS, OTHERWISE IT ALWAYS MOVES IN THE OPPOSITE DIRECTION OF THE STATE.

principal's favorite level $\gamma \star f$ almost surely, therefore the distortion accumulates over time on average. It is shown in the proof that the continuation cost C_t always has a non-negative drift, strictly positive if the target is not a constant.

Part (ii) predicts the drift of the continuation payoff in two cases, based on the difference $\gamma \star f - f$, which can be equivalently expressed as $\frac{1}{2r}\big((\gamma \star f)', (\gamma \star f)''\big) \cdot (2\mu, \sigma^2)$. In the case of $\mu = 0$, the difference is solely determined by the curvature of $\gamma \star f$. This result implies that the agent does not necessarily end up immiserated; instead, his destiny depends on the nature of the target function.

D. Examples

The specific form of the target function varies by economic situations. This subsection studies some typical target functions and their implications for the optimal contract.

Binary Target

In some applications, the target takes only binary values. For instance, the market size enjoys a discrete boost if a product-related parameter satisfies a minimum requirement, say, zero. In other words, there is an "active zone" $[0, \infty)$ in which the target is higher.

To be specific, let $f(\theta) = \mathbb{1}\{\theta \geqslant 0\}$. While the target is discontinuous at 0, the expected future target $\gamma \star f$ is continuously differentiable. It can be shown that

$f' < (\gamma \star f)'$ at all states except zero, the discontinuity point. Since the target is flat at these states, the action is neither conformist nor contrarian. However, the action *does* respond to the state in the opposite direction. Only at the cutoff state 0 does the action jump in the same direction and with the same magnitude as the target. Panel (a) of Figure 3 pictures the target and the expected future target. Panel (b) simulates the time paths for the target and the action. Every time the target jumps, the optimal action jumps one-for-one. This happens when the state crosses the cutoff 0. At other times, the action moves against the state, which is not shown in the picture.

The intuition is as follows. When the state increases from 1 to 2, the current target does not change (stays at 1), but on average the future target increases. As a result, the allocation should favor the future by lowering the action now to make room for future increases.

Kinked Target

In some situations, the target action has a piece-wise nature. Initially the target increases fast with the state, until the state reaches some threshold beyond which the target becomes less responsive, or even stops growing. Such regime changes can occur when some sub-market is saturated, or when the marginal return for a technical parameter drops after meeting a certain threshold. Of course, the regime shift can occur at multiple states with any type of kinks.

For concreteness, I study an example where the target function increases one-for-one at states below 0 but the slope drops to $b \in (0,1)$ at states above 0:

$$f(\theta) = \begin{cases} \theta & \text{if } \theta \leqslant 0 \\ b\theta & \text{if } \theta > 0 \end{cases} .$$

A little algebra shows that $f'(\theta) < (\gamma \star f)'(\theta)$ if and only if $\theta > 0$, which, based on Theorem 3, predicts conformist actions at states below 0 but contrarian actions above 0. The target function and its expected future version are depicted in Panel (a) of Figure 4. Panel (b) shows the simulated paths for the target and the action. Perfectly aligned with the analysis, the action co-moves with the target whenever the target $f(\theta)$ is below the kink ($\theta < 0$), and it moves against the target above the kink.

Intuitively, at low states the information sensitivity is already at its highest

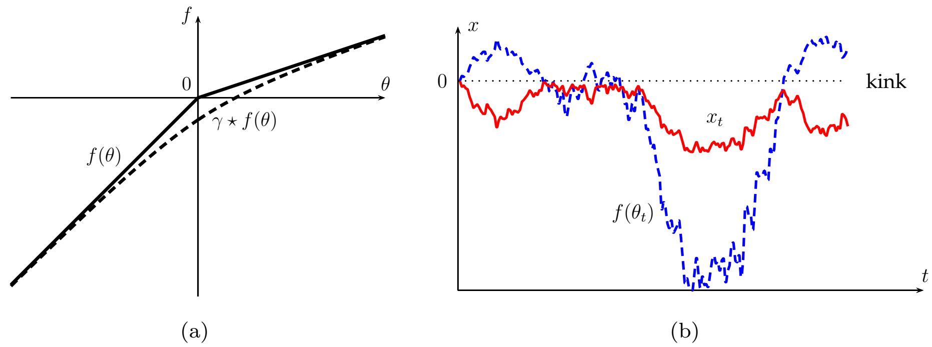

FIGURE 4. KINKED TARGET WITH PARAMETERS: $r = 1, \sigma = 1, \mu = 0, \theta_0 = 0, b = 0.33$. PANEL (A): THE TARGET FUNCTION (SOLID CURVE) AND ITS CONVOLUTED VERSION (DASHED CURVE). PANEL (B): SIMULATED PATHS OF THE TARGET (BLUE DASHED CURVE) AND THE OPTIMAL ACTION (RED SOLID CURVE). THE DOTTED CURVE SHOWS THE LOCUS OF THE TARGET AT THE KINK. THE CYCLICALITY OF THE TWO PATHS DEPENDS ON WHETHER THE TARGET IS ABOVE THE KINK.

possible value, and thus the expected future slope can only be lower. At high states where the information sensitivity is already lowest, it can only increase in the future. This explains the conformist and contrarian patterns.

Exponential Target

Exponential target can be a good approximation if the target is monotone in the state but displays increasing or decreasing sensitivity to state changes. For instance, suppose the state is the profitability of a product in an industry where an increase in profitability calls for disproportional increase in the level of operation.

Generally, an exponential target function can be written as $f(\theta) = b_0 e^{b_1 \theta}$ such that Assumption 1 is satisfied. The target is increasing if $b_0 b_1 > 0$ and decreasing if $b_0 b_1 < 0$. With this exponential target function, the expected future marginal target becomes $(\gamma \star f)'(\theta) = \frac{2r}{2r - 2\mu b_1 - \sigma^2 b_1^2} f'(\theta)$.

According to Theorem 3, the action is contrarian at *all* states as long as b_1 is not between 0 and $-\frac{2\mu}{\sigma^2}$. When $\mu = 0$, this condition is automatically satisfied. Panel (a) of Figure 5 shows both the target function and the expected future target function, when the target function is chosen to be exponential and the drift is zero. Panel (b) displays simulated paths of the target and the optimal action. The dotted curve represents the *same* trend for both paths (recall that the action is "correct on average"). The action path and the target path, plotted

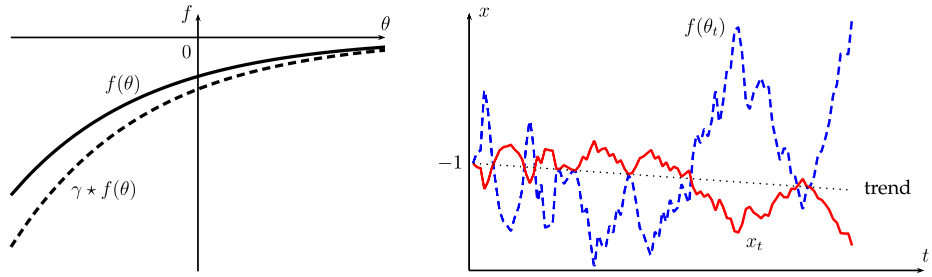

(a) (b)

FIGURE 5. EXPONENTIAL TARGET WITH PARAMETERS: $r = 1, \sigma = 1, \mu = 0, \theta_0 = 0, b_0 = -1, b_1 = -0.7$. PANEL (A): THE TARGET FUNCTION (SOLID CURVE) AND ITS CONVOLUTED VERSION (DASHED CURVE). PANEL (B): SIMULATED PATHS OF THE TARGET (BLUE DASHED CURVE) AND THE OPTIMAL ACTION (RED SOLID CURVE). THE DOTTED CURVE IS THE SAME TREND FOR BOTH PATHS; IT ALWAYS SEPARATES THE TARGET AND THE ACTION ON OPPOSITE SIDES.

in solid and dashed curves respectively, always lie on opposite sides of the dotted curve.

It is important to compare this example with the previous one of kinked target. Both represent an increasing and concave target, but the implications for the optimal contract are very different. This contrast corroborates the earlier finding that it is insufficient to only look at the curvature of the target function to determine the pattern of the optimal contract.

V. Extensions

This section extends the main model in two directions: to consider less persistent state process by introducing mean reversion, and to partially relax the no-transfer assumption by allowing for money with limited liability.

A. Mean Reverting State Process

Persistence of the state process has demonstrated its importance in the intertemporal trade-off: future information sensitivity can be important, and sometimes even more important than the current information sensitivity because of the persistence in the state. In the main model the persistence is high in that any current shock persists through time without decay.

When the state process exhibits mean-reversion, the persistence is weaker. For

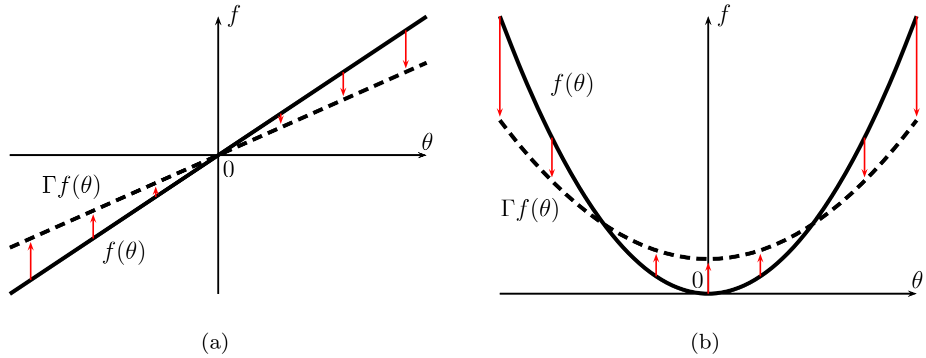

FIGURE 6. THE EFFECT OF MEAN-REVERSION, WITH PARAMETERS $r = 1, \sigma = 1, \phi = .5, \theta_0 = 0$. PANEL (A): THE SOLID CURVE IS THE TARGET $f(\theta) = \theta$ AND THE DASHED CURVE IS Γf. PANEL (B): THE SOLID CURVE IS THE TARGET $f(\theta) = \theta^2$ AND THE DASHED CURVE IS Γf.

simplicity, here I assume a mean-reverting state process:

$$\mathrm{d}\theta_t = -\phi(\theta_t - \theta_0)\mathrm{d}t + \sigma \mathrm{d}Z_t.$$

It can be shown that Proposition 1 still holds as a local version. With the same procedure as in the main model, the cost and policy functions are obtained as follows:

$$C(\theta, W) = (W - \Gamma f(\theta))^2 + \frac{\sigma^2}{r}\Gamma(\Gamma f)'^2(\theta),\ x(\theta, W) = W + f(\theta) - \Gamma f(\theta),$$

where the functional operator Γ is such that $g = \Gamma f$ gives the unique solution to the ODE:

$$\sigma^2 g''(\theta) - 2(\theta - \theta_0)\mu g'(\theta) - 2rg(\theta) = -2rf(\theta), \quad \lim_{\theta \to \pm\infty} e^{-\sqrt{\frac{r}{2\sigma^2}}|\theta|} g(\theta) = 0. \tag{14}$$

Unlike the main model, the explicit form of Γ is not obtainable in general with mean reversion. To gain an impression of the effect of mean-reversion, it nonetheless suffices to look at some specific cases where closed-form solution is available.

To begin with, let $\theta_0 = 0$ and $f(\theta) = \theta$, which is linear in the state. By method of undetermined coefficients, it is easy to verify the solution $\Gamma f(\theta) = \frac{r}{r+\phi}\theta$. While the solution is still linear, the coefficient on θ is *dampened towards zero* by a factor of $\frac{r}{r+\phi} < 1$. Intuitively, states far away from θ_0 are very likely to drift back towards θ_0 and hence take less weight than states near θ_0 in the computation of the expected future target. As a result, future information sensitivity is dampened. Recall that communication is babbling and the action does not respond to information when $\phi = 0$, this dampening effect causes two changes: communication becomes effective and the action is always conformist. As $\mu \to \infty$, the state process approaches i.i.d., and Γf is completely flattened. In that case, the action tracks the target exactly one-for-one, and the complete information cost obtains. This example also explains the prevalence of "conformist" quota usage in the literature where the private information follows a finite state Markov chain and the target is linear. Panel (a) of Figure 6 how the future information sensitivity $(\Gamma f)'$ is dampened towards zero relative to the current information sensitivity f'.

Next, let $f(\theta) = \theta^2$ be quadratic. The solution becomes $\Gamma f(\theta) = \frac{\sigma^2}{r+2\phi} + \frac{r}{r+2\phi}\theta^2$. This time, the coefficient for θ^2 is dampened even more by $\frac{r}{r+2\phi}$. Communication is effective, and the action is conformist as long as $f' \neq 0$. When $\phi = 0$, we are back in the main model: communication fails and the action path does not respond to information. Panel (b) of Figure 6 belongs to this quadratic case.

If we take a step further to let $f(\theta) = \theta^3$, the solution $\Gamma f = \frac{r}{r+3\phi}\theta^3 + \frac{3r\sigma^2}{(r+3\phi)(r+\phi)}\theta$ is more interesting. The action would have been contrarian at all nonzero states if $\phi = 0$, but now with mean reversion the set of states shrinks to $\left(-\sqrt{\frac{r}{3\phi(r+\phi)}}\sigma, \sqrt{\frac{r}{3\phi(r+\phi)}}\sigma\right) \setminus \{0\}$. As $\phi \to \infty$, the measure of this set vanishes.

In sum, the mean reversion serves to undermine the importance of future information sensitivity, resulting in a more conformist action than before.

B. *Transfer with Limited Liability*

There are situations where monetary transfer is either legal or difficult to detect or prohibit. How do the main results extend? If the transfer is entirely unconstrained, then the socially efficient action is always taken, and the principal always uses (positive or negative) monetary transfer to off set the effect of action on the agent.

More realistically, money only moves from the principal to the agent, i.e., the agent has a limited liability constraint. For simplicity I focus on the linear target function: $f(\theta) = \theta$. The linearity would have resulted in babbling were there no transfer allowed, but it is no longer the case here. Let $y = (y_t \geqslant 0)_{t \geqslant 0}$ denote the process of non-negative transfer from the principal to the agent, as part of the contract. Both players are risk-neutral with respect to money. The principal solves:

$$\min_{\substack{(x_t(\cdot))_{t\geqslant 0}\\(y_t(\cdot)\geqslant 0)_{t\geqslant 0}}} \quad \mathbb{E}\left[\int_0^\infty re^{-rt}\left((x_t(\theta^t)-f(\theta_t))^2+y_t(\theta^t)\right)\mathrm{d}t\right]$$

$$\text{s.t.}\quad \mathbb{E}\left[\int_0^\infty re^{-rt}(x_t(\theta^t)+y_t(\theta^t))\mathrm{d}t\right]\geqslant \mathbb{E}\left[\int_0^\infty re^{-rt}(x_t(\hat{\theta}^t)+y_t(\hat{\theta}^t))\mathrm{d}t\right],$$

$$\text{where } \hat{\theta}_t=\theta_t+\int_0^t m_s\mathrm{d}s,\ \forall\, m\in\mathcal{M}.$$

It can be shown that in the optimal contract, transfer never occurs in finite time. Suppose it is used at some time t, then there is always another contract delaying this payment with interest rate r that keeps the incentive of the agent but relaxes the limited liability. However, the optimal contract is affected by the mere existence of money. In fact, money serves as an option that can be used to fulfill the continuation payoff W. If W is excessively high relative to the current target, then the principal is promising too much, and has the cheaper option to use money rather than high actions to fulfill the promise, considering the fact that money has a linear instead of quadratic cost. In the limit as $W\to\infty$, the cost function converges to the "unlimited money" case. In contrast, when W is excessively low, the principal wants to "charge" money from the agent, but cannot due to limited liability. Hence, as $W\to-\infty$, the cost function converges to the no-transfer case as in the main model.

Since $f(\theta)=\theta$ is homogeneous, the cost function depends only on the difference $W-\theta$: $C(\theta,W)=C(W-\theta)$, with abuse of notations. The uni-variate function C is shown in Panel (a) of Figure 7. As expected, it approaches the unlimited-money lower bound as $W-\theta\to\infty$, and the no-money upper bound as $W-\theta\to-\infty$. The policy function reads $x(\theta,W)=\theta+\frac{1}{2}C'(W-\theta)$. This time, $\frac{\partial x(\theta,W)}{\partial\theta}$ depends on W, meaning that whether the action is conformist or contrarian at a given state is history-dependent. Panel (b) of Figure 7 represents a ratio of derivatives $\frac{\partial x(\theta,W)}{\partial\theta}/\frac{\mathrm{d}f(\theta)}{\mathrm{d}\theta}$ fixing $W=0$. When θ is high so that W is relatively low, this ratio is nearly zero, consistent with the non-responsiveness result in the no-money case when f is linear. When θ is very negative, the derivative is close to one, consistent with the perfect conformist action in the unlimited-money case.

Similar results extend to nonlinear target functions. Panel (c) of Figure 7 shows

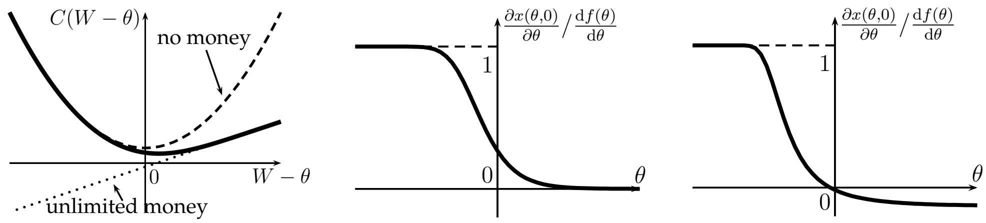

(a) (b) (c)

FIGURE 7. TRANSFER WITH LIMITED LIABILITY. PARAMETERS: $r = 1, \sigma = 1$. PANEL (A): COST FUNCTION IN SOLID CURVE, WITH $f(\theta) = \theta$. PANEL (B): THE RATIO OF DERIVATIVES AT $W = 0$ IN SOLID CURVE, WITH $f(\theta) = \theta$. PANEL (C): THE RATIO OF DERIVATIVES AT $W = 0$ IN SOLID CURVE, WITH $f(\theta) = -e^{-0.6\theta}$.

the ratio of derivatives holding $W = 0$, for the target function $f(\theta) = -e^{-0.6\theta}$. Again, its behavior approaches either contrarian or perfect conformist based on the tightness of money.

VI. Conclusion

I use the principal-agent model with dynamic contract to study the communication problem. Since the agent has state-independent preferences over the principal's actions, one-shot communication is inevitably babbling even if the principal can commit. In contrast, I show that a dynamic contract salvages partial value of information in most cases, because of the principal's ability to reallocate distortions across time while respecting the incentives of the agent. Nonetheless, there are cases in which the ability of inter-temporal trade-off disappears even if the principal's favorite action is non-linear the state.

More importantly, the optimal contract can behave in a counter-intuitive manner: decrease the action when the target increase, and vice versa, despite the obvious temptation to close the gap between the action and the target. This phenomenon arises when future weighs more than present in terms of sensitivity to information. I show that it is the third derivative, not only the curvature, of the target function that plays a critical role in determining the shape of the contract.

The contrarian response can be viewed as a new implication from agency problems; it never arises if there is no conflict of interests. The more aligned the preferences, the less likely that the optimal contract is contrarian. The contrarian action against the agent's report does not come from the distrust towards the agent (recall that it is a truthful contract), instead it can be the most efficient way of using information given the conflict of interests.

REFERENCES

Abreu, Dilip, David Pearce, and Ennio Stacchetti. 1986. "Optimal cartel equilibria with imperfect monitoring." *Journal of Economic Theory*, 39(1): 251–269.

Alonso, Ricardo, and Niko Matouschek. 2008. "Optimal delegation." *The Review of Economic Studies*, 75(1): 259–293.

Amador, Manuel, and Kyle Bagwell. 2013. "The theory of optimal delegation with an application to tariff caps." *Econometrica*, 81(4): 1541–1599.

Antič, Nemanja, and Kai Steverson. 2016. "Screening through coordination." *Working paper.*

Aumann, Robert J., and Sergiu Hart. 2003. "Long cheap talk." *Econometrica*, 71(6): 1619–1660.

Battaglini, Marco. 2002. "Multiple referrals and multidimensional cheap talk." *Econometrica*, 70(4): 1379–1401.

Battaglini, Marco. 2005. "Long-term contracting with Markovian consumers." *American Economic Review*, 95(3): 637–658.

Boleslavsky, Raphael, and Tracy R Lewis. 2016. "Evolving influence: Mitigating extreme conflicts of interest in advisory relationships." *Games and Economic Behavior*, 98: 110–134.

Chakraborty, Archishman, and Rick Harbaugh. 2010. "Persuasion by cheap talk." *The American Economic Review*, 100(5): 2361–2382.

Crawford, Vincent P., and Joel Sobel. 1982. "Strategic information transmission." *Econometrica*, 50(6): 1431–1451.

DeMarzo, Peter M, and Yuliy Sannikov. 2016. "Learning, termination, and payout policy in dynamic incentive contracts." *The Review of Economic Studies*, 84(1): 182–236.

Fernandes, Ana, and Christopher Phelan. 2000. "A recursive formulation for repeated agency with history dependence." *Journal of Economic Theory*, 91(2): 223–247.

Fershtman, Daniel. 2017. "Dynamic delegation: Specialization and favoritism." *Working paper.*

Golosov, Mikhail, Vasiliki Skreta, Aleh Tsyvinski, and Andrea Wilson. 2014. "Dynamic strategic information transmission." *Journal of Economic Theory*, 151: 304–341.

Goltsman, Maria, Johannes Hörner, Gregory Pavlov, and Francesco Squintani. 2009. "Mediation, arbitration and negotiation." *Journal of Economic Theory*, 144(4): 1397 – 1420.

Green, Jerry R., and Nancy L. Stokey. 2007. "A two-person game of information transmission." *Journal of Economic Theory*, 135(1): 90 – 104.

Guo, Yingni, and Johannes Hörner. 2017. "Dynamic mechanisms without money." *Working paper.*

He, Zhiguo, Bin Wei, Jianfeng Yu, and Feng Gao. 2017. "Optimal long-term contracting with learning." *The Review of Financial Studies*, 30(6): 2006–2065.

Holmström, Bengt Robert. 1977. "On incentives and control in organizations."

Jackson, Matthew O, and Hugo F Sonnenschein. 2007. "Overcoming incentive constraints by linking decisions." *Econometrica*, 75(1): 241–257.

Kapička, Marek. 2013. "Efficient allocations in dynamic private information economies with persistent shocks: a first-order approach." *Review of Economic Studies*, 80(3): 1027–1054.

Krishna, Vijay, and John Morgan. 2001. "A model of expertise." *The Quarterly Journal of Economics*, 116(2): 747–775.

Krishna, Vijay, and John Morgan. 2004. "The art of conversation: eliciting information from experts through multi-stage communication." *Journal of Economic Theory*, 117(2): 147 – 179.

Malenko, Andrey. 2016. "Optimal dynamic capital budgeting." *Working paper.*

Margaria, Chiara, and Alex Smolin. 2017. "Dynamic communication with biased senders." *Games and Economic Behavior.*

Pavan, Alessandro, Ilya Segal, and Juuso Toikka. 2014. "Dynamic mechanism design: A myersonian approach." *Econometrica*, 82(2): 601–653.

Renault, Jérôme, Eilon Solan, and Nicolas Vieille. 2013. "Dynamic sender–receiver games." *Journal of Economic Theory*, 148(2): 502–534.

Sannikov, Yuliy. 2008. "A continuous-time version of the principal-agent problem." *The Review of Economic Studies*, 75(3): 957–984.

Thomas, Jonathan, and Tim Worrall. 1990. "Income fluctuation and asymmetric information: An example of a repeated principal-agent problem." *Journal of Economic Theory*, 51(2): 367–390.

Williams, Noah. 2011. "Persistent private information." *Econometrica*, 79(4): 1233–1275.

Appendix

Solving the Two-Period Contract

PROOF:

The IC's are simplified to one equation: $x_1(\theta_1) + x_2(\theta_1) = W$. Plug this back to the objective to obtain the unconstrained problem:

$$\min_{x_1(\cdot),W} \mathbb{E}\left[(x_1(\theta_1) - f(\theta_1))^2 + \mathbb{E}\left[(W - x_1(\theta_1) - f(\theta_2))^2 \,|\theta_1\right]\right].$$

For every θ_1, the FOC w.r.t. $x_1(\theta_1)$ gives

$$x_1(\theta_1) = \frac{1}{2}W + \frac{1}{2}\left(f(\theta_1) - \mathbb{E}\left[f(\theta_2)|\theta_1\right]\right).$$

Plug the above to the objective, and then take FOC w.r.t. W:

$$W = \mathbb{E}f(\theta_1) + \mathbb{E}f(\theta_2).$$

Replacing W in the expression of $x_1(\theta_1)$ with the above, we have the solution (3). The IC condition then leads to (4).

PROOF OF LEMMA 1

PROOF:

Suppose the given contract x induces a (not necessarily truthful) strategy $m \in \mathcal{M}$, which generates a mapping from state paths into action paths. Let $M_t \equiv \int_0^t m_s \mathrm{d}s$ be the accumulated manipulation. Consider a new contract $x^\dagger$ such that $x_t^\dagger(\hat{\theta}^t) \equiv x_t((\hat{\theta} + M)^t)$. I claim that truth-telling $m^\dagger$ is optimal for the agent under the new contract. If not, then there exists a strategy $m' \in \mathcal{M}$ along with $M'_t \equiv \int_0^t m'_s \mathrm{d}s$ such that $\mathbb{E}\left[\int_0^\infty re^{-rt} x_t^\dagger((\theta + M')^t)\mathrm{d}t\right] > \mathbb{E}\left[\int_0^\infty re^{-rt} x_t^\dagger(\theta^t)\mathrm{d}t\right]$. Contradiction arises as $m + m' \in \mathcal{M}$ outperforms m in the original contract:

$$
\begin{aligned}
&\mathbb{E}\left[\int_0^\infty re^{-rt} x_t((\theta + M + M')^t)\mathrm{d}t\right] = \mathbb{E}\left[\int_0^\infty re^{-rt} x_t^\dagger((\theta + M')^t)\mathrm{d}t\right] \\
> \;&\mathbb{E}\left[\int_0^\infty re^{-rt} x_t^\dagger(\theta^t)\mathrm{d}t\right] = \mathbb{E}\left[\int_0^\infty re^{-rt} x_t((\theta + M)^t)\mathrm{d}t\right].
\end{aligned}
$$

The new contract $x^\dagger$ implements the original mapping from θ^t to x^t by construction, $\forall\, t$.

PROOF OF LEMMA 2

PROOF:

Given a contract x, define the process of the agent's total payoff evaluated at time 0 but with information at time t:

$$
\hat{W}_t^0 \equiv \int_0^t re^{-rs} x_s \mathrm{d}s + e^{-rt} W_t,
$$

which is a martingale because for any $0 \leqslant t' \leqslant t$,

$$
\mathbb{E}_{t'} \hat{W}_t^0 = \int_0^{t'} re^{-rs} x_s \mathrm{d}s + \mathbb{E}_{t'}\left[\int_{t'}^t re^{-rs} x_s \mathrm{d}s\right] + e^{-rt}\mathbb{E}_{t'}\left[\int_t^\infty re^{-r(s-t)} x_s \mathrm{d}s\right] = \hat{W}_{t'}^0.
$$

By Theorem 1.3.13 in Karatzas and Shreve (1991), the martingale $\hat{W}_t^0$ has a RCLL modification. Therefore by Theorem 3.4.15 in the same book, the martingale has a representation

$$\hat{W}_t^0 = \hat{W}_0^0 + \int_0^t re^{-rs}\beta_s\sigma\mathrm{d}Z_s,\ \forall\ t \geqslant 0.$$

Subtracting the two expressions for $\hat{W}_t^0$ and then differentiating w.r.t. t, we have

$$\mathrm{d}W_t = r(W_t - x_t)\mathrm{d}t + r\beta_t\sigma\mathrm{d}Z_t = r(W_t - x_t)\mathrm{d}t + r\beta_t(\mathrm{d}\hat{\theta}_t - \mu\mathrm{d}t),$$

which has an equivalent integral form $W_t = W_0 + \int_0^t r(W_s - x_s)\mathrm{d}s + \int_0^t r\beta_s\sigma\mathrm{d}Z_s$.

Proof of Proposition 1

PROOF:

For any strategy $m \in \mathcal{M}$, Novikov's condition is satisfied. By Girsanov Theorem there exists a martingale Y with $Y_t \equiv e^{\frac{1}{\sigma}\int_0^t m_s\mathrm{d}Z_s - \frac{1}{2\sigma^2}\int_0^t m_s^2\mathrm{d}s}$, serving as the Radon-Nikodym derivative between the measure induced by m and the measure under truth-telling. It evolves according to $\mathrm{d}Y_t = Y_t\frac{m_t}{\sigma}\mathrm{d}Z_t$ with $Y_0 = 1$. Besides Y_t, the cumulative manipulation $M_t = \int_0^t m_s\mathrm{d}s$ is also a state variable, with evolution $\mathrm{d}M_t = m_t\mathrm{d}t$. Then, the agent's payoff from a strategy $m \in \mathcal{M}$ is $\mathbb{E}\left[\int_0^\infty re^{-rt}Y_t x_t\mathrm{d}t\right]$.

Let p^Y be the costate variable for the drift of Y, and q^Y the costate for the volatility of Y. Let p^M and q^M be the counterparts for M. The agent's current value Hamiltonian is $rYx + q^Y\frac{Ym}{\sigma} + p^M m$.

The first order condition for $m = 0$ to be optimal, evaluated at $m = 0, Y = 1$, is

(A1) $$\frac{q^Y}{\sigma} + p^M = 0.$$

The Euler equations for Y and M, evaluated at $m = 0, Y = 1$, are

(A2) $$\mathrm{d}p^Y = r(p^Y - x)\mathrm{d}t + \frac{q^Y}{\sigma}(\sigma\mathrm{d}Z_t),$$

(A3) $$\mathrm{d}p^M = rp^M\mathrm{d}t + \frac{q^Y}{\sigma}(\sigma\mathrm{d}Z_t),$$

with transversality conditions $\lim_{t\to\infty} p_t^Y e^{-rt} = 0$ and $\lim_{t\to\infty} p_t^M e^{-rt} = 0$. The solution to the above BSDE's are $p_t^Y = \mathbb{E}_t\left[\int_t^\infty re^{-r(s-t)}x_s\mathrm{d}s\right] = W_t$ and $p_t^M = 0$, where W_t is the agent's continuation payoff defined in Section III.B. Hence, by comparing (A2) and (8), we have $\frac{q^Y}{\sigma} = r\beta$. Plugging this back to (A1) and using the fact $p^M = 0$, we have the necessary condition $\beta = 0$.

PROOF OF THEOREM 1

PROOF:

The proof takes two steps. Step one, I show that the candidate cost function, along with policy function (12), indeed achieves the lowest cost in the relaxed problem. Step two, I show that the candidate solution is also satisfies the global IC conditions.

Step one. For any contract $\hat{x}$ satisfying the IC necessary condition and the transversality condition of the agent, define $\hat{W}$ as the resulting continuation payoff process with $\hat{W}_0 = W_0$, and define

$$\hat{C}_t^0 \equiv \int_0^t re^{-rs}(\hat{x}_s - f(\theta_s))^2 \mathrm{d}s + e^{-rt}C^*(\theta_t, \hat{W}_t) \tag{A4}$$

as the total cost process evaluated at time t. In this process, the policy follows the arbitrary contract $\hat{x}$ until time t and then the candidate cost function takes place as continuation, promising $\hat{W}_t$ as continuation payoff. The goal is to show that $\hat{C}_t^0$ is a martingale if $\hat{x}$ coincides with the optimal policy (12), and is a sub-martingale if not. The total differential for $\hat{C}_t^0$ is

$$\begin{aligned} e^{rt}\mathrm{d}\hat{C}_t^0 &= r(\hat{x}_t - f(\theta_t))^2\mathrm{d}t - rC^*(\theta_t, \hat{W}_t)\mathrm{d}t + r(\hat{W}_t - \hat{x}_t)C_W^*(\theta_t, \hat{W}_t)\mathrm{d}t \\ &\quad + \sigma C_\theta^*(\theta_t, \hat{W}_t)\mathrm{d}Z_t + \mu C_\theta^*(\theta_t, \hat{W}_t)\mathrm{d}t + \frac{\sigma^2}{2}C_{\theta\theta}^*(\theta_t, \hat{W}_t)\mathrm{d}t \\ &= \sigma C_\theta^*(\theta_t, \hat{W}_t)\mathrm{d}Z_t + r(\hat{x}_t - x_t)(\hat{x}_t + x_t - 2f(\theta_t) - C_W^*(\theta_t, \hat{W}_t))\mathrm{d}t \\ &= \sigma C_\theta^*(\theta_t, \hat{W}_t)\mathrm{d}Z_t + r(\hat{x}_t - x_t)^2\mathrm{d}t, \end{aligned}$$

where $x_t = x^*(\theta_t, \hat{W}_t)$ is the candidate policy, the second equality follows from the HJB (9), and the third equality utilizes the policy function (12). It is clear that

$$e^{rt}\frac{\mathbb{E}_t\mathrm{d}\hat{C}_t^0}{\mathrm{d}t} = r(\hat{x}_t - x_t)^2 \geqslant 0, \tag{A5}$$

with equality if and only if $\hat{x}_t = x_t$.

In the following I show that the arbitrary contract $\hat{x}$ does not achieve a lower cost, given the agent's truthful report. For any initial value (θ_0, W_0),

$$C^*(\theta_0, W_0) = \hat{C}_0^0 \leqslant \mathbb{E}\hat{C}_\infty^0 = \mathbb{E}\int_0^\infty re^{-rs}(\hat{x}_s - f(\theta_s))^2 \mathrm{d}s + \lim_{t\to\infty} \mathbb{E}e^{-rt}C^*(\theta_t, \hat{W}_t).$$

If $\mathbb{E}\int_0^\infty re^{-rs}(\hat{x}_s - f(\theta_s))^2 \mathrm{d}s = \infty$, then this contract $\hat{x}$ results in infinite cost. Now suppose $\mathbb{E}\int_0^\infty re^{-rs}(\hat{x}_s - f(\theta_s))^2 \mathrm{d}s < \infty$, i.e. $\mathbb{E}\int_0^\infty (\hat{x}_s - f(\theta_s))^2 \mathrm{d}\left(-e^{-rs}\right) < \infty$. This means, with respect to the finite product measure, $\hat{x} - f(\theta) \in \mathcal{L}^2$. At the same time, it is straightforward to verify that $f(\theta) \in \mathcal{L}^2$ with Assumption 1 (ii). By the closure to addition of $\mathcal{L}^2$, one arrives at the conclusion that $\mathbb{E}\int_0^\infty re^{-rs}\hat{x}_s^2 \mathrm{d}s < \infty$. For any $\hat{W}$ satisfying the agent's transversality condition, we must have $\hat{W}_t = \mathbb{E}_t \int_0^\infty \hat{x}_{t+s}\mathrm{d}\left(-e^{-rs}\right)$. Therefore,

$$\begin{aligned} e^{rt}\mathbb{E}\hat{W}_t^2 &= e^{rt}\mathbb{E}\left(\int_0^\infty \hat{x}_{t+s}\mathrm{d}\left(-e^{-rs}\right)\right)^2 \\ &\leqslant e^{rt}\mathbb{E}\int_0^\infty \hat{x}_{t+s}^2 \mathrm{d}\left(-e^{-rs}\right) = \mathbb{E}\int_t^\infty re^{-rs}\hat{x}_s^2 \mathrm{d}s < \infty, \end{aligned}$$

where the first inequality follows from Hölder's inequality. Taking limit as $t \to \infty$, we have $\lim_{t\to\infty} \mathbb{E}\left[e^{rt}\hat{W}_t^2\right] = 0$. Moreover, it is straightforward to verify that $\mathbb{E}\left[e^{rt}(\gamma \star f(\theta_t))^2\right]$ vanishes. Hence, $\lim_{t\to\infty} \mathbb{E}\left[e^{-rt}C^*(\theta_t, \hat{W}_t)\right] = 0$ by noticing that $(\hat{W}_t - \gamma \star f(\theta_t))^2 \leqslant 2(\hat{W}_t^2 + (\gamma \star f(\theta_t))^2)$. This means $C^*(\theta_0, W_0) \leqslant \mathbb{E}\int_0^\infty re^{-rs}(\hat{x}_s - f(\theta_s))^2 \mathrm{d}s$.

Step two. It remains to check the IC conditions for global deviations. Suppose the agent adopts an arbitrary misreporting strategy m, so that the reported process is $\hat{\theta} = \theta + M$ where $M_t = \int_0^t m_t \mathrm{d}t$. The resulting action and continuation payoff processes are denoted as x^m and W^m. Since $\mathrm{d}W_t^m = r(W_t^m - x_t^m)\mathrm{d}t$, we know that $x_t^m \mathrm{d}t = -\frac{e^{rt}}{r}\mathrm{d}\left(e^{-rt}W_t^m\right)$. Therefore the agent's payoff from strategy m is:

$$\lim_{t\to\infty} \mathbb{E}\int_0^t re^{-rs}x_s^m \mathrm{d}s = W_0 - \lim_{t\to\infty} e^{-rt}\mathbb{E}W_t^m,$$

which means the IC conditions are met as long as $\lim_{t\to\infty} e^{-rt}\mathbb{E}W_t^m = 0$ for all $m \in \mathcal{M}$. For any given strategy $m \in \mathcal{M}$, there exists $T > 0$ such that

$\mathbb{E}e^{\frac{\sqrt{2r\sigma^2+\mu^2}}{\sigma^2}M_t} < e^{rt}$ for all $t > T$, so for those large t,

$$\begin{aligned} &\int_0^t \mathbb{E}e^{\alpha_1(\theta_s+M_s)}\mathrm{d}s \leqslant \int_0^t \sqrt{\mathbb{E}e^{2\alpha_1\theta_s}}\sqrt{\mathbb{E}e^{2\alpha_1 M_s}}\mathrm{d}s \\ \leqslant\ & e^{\alpha_1\theta_0}\int_0^t e^{\alpha_1(\mu+\alpha_1\sigma^2)s}\left(\mathbb{E}e^{2\overline{\alpha}M_s}\right)^{\frac{\alpha_1}{2\overline{\alpha}}}\mathrm{d}s \\ =\ & e^{\alpha_1\theta_0}\int_0^T e^{\alpha_1(\mu+\alpha_1\sigma^2)s}\left(\mathbb{E}e^{2\overline{\alpha}M_s}\right)^{\frac{\alpha_1}{2\overline{\alpha}}}\mathrm{d}s + e^{\alpha_1\theta_0}\int_T^t e^{\alpha_1(\mu+\alpha_1\sigma^2)s}\left(\mathbb{E}e^{2\overline{\alpha}M_s}\right)^{\frac{\alpha_1}{2\overline{\alpha}}}\mathrm{d}s \\ \leqslant\ & e^{\alpha_1\theta_0}\int_0^T e^{\alpha_1(\mu+\alpha_1\sigma^2)s}\left(\mathbb{E}e^{2\overline{\alpha}M_s}\right)^{\frac{\alpha_1}{2\overline{\alpha}}}\mathrm{d}s + \frac{e^{\alpha_1\theta_0}}{\alpha_2}\left(e^{\alpha_2 t} - e^{\alpha_2 T}\right), \end{aligned}$$

where $\alpha_2 \equiv \alpha_1(\mu+\alpha_1\sigma^2) + r\frac{\alpha_1}{2\overline{\alpha}} < r$. Hence, the first term in the last line is finite while the second term grows slower than e^{rt}. Similarly, $\int_0^t \mathbb{E}e^{-\alpha_1(\theta_s+M_s)}\mathrm{d}s$ grows slower than e^{rt} too. With the candidate policy function,

$$\begin{aligned} &\left|\frac{\mathrm{d}W_t^m}{\mathrm{d}t}\right| = r\left|\gamma\star f(\theta_t+M_t) - f(\theta_t+M_t)\right| \\ \leqslant\ & r\alpha_0\frac{2\alpha_1\mu+4r-\alpha_1^2\sigma^2}{2\alpha_1\mu+2r-\alpha_1^2\sigma^2}e^{-\alpha_1(\theta_t+M_t)} + r\alpha_0\frac{-2\alpha_1\mu+4r-\alpha_1^2\sigma^2}{-2\alpha_1\mu+2r-\alpha_1^2\sigma^2}e^{\alpha_1(\theta_t+M_t)}, \end{aligned}$$

therefore with the above analysis, $\lim_{t\to\infty} e^{-rt}\mathbb{E}W_t^m = \lim_{t\to\infty} e^{-rt}\mathbb{E}|W_t^m| = 0$.

Finally, to obtain the lowest cost for the principal as well as the explicit optimal contract, I set W_0 optimally at $\gamma\star f(\theta_0)$. Also, with the policy function the continuation payoff evolves as:

$$W_t = W_0 + \int_0^t \mathrm{d}W_s = \gamma\star f(\theta_0) + r\int_0^t(\gamma\star f(\theta_s) - f(\theta_s))\mathrm{d}s.$$

Hence, plug W_t in (12) to obtain (13).

Proof for Theorem 2

PROOF:

For part (i), the "if" direction is trivial. For the "only if" direction, since $(\gamma\star f)'^2 \geqslant 0$ and the convolution preserves the sign, we know that $\gamma\star((\gamma\star f)'^2)(\theta_0) = 0$ if and only if $(\gamma\star f)' \equiv 0$ almost everywhere. By Assumption 1, $\gamma\star f$ is continuously differentiable, so $\gamma\star f$ is a constant, which means f is a constant almost everywhere.

For part (ii), the "if" direction can be verified by plugging $f(\theta) = c_0 + c_1\theta + c_2\theta^2$. The resulting action path $x_t = f(\theta_0) - \frac{\sigma^2}{2}2c_2t$ is deterministic, achievable in a babbling equilibrium. For the "only if" direction, define $\hat{C}_t^0$ the same way as in (A4) for the babbling contract $\hat{x}$. The drift of $\hat{C}_t^0$ satisfies $e^{rt}\frac{\mathbb{E}_t[\mathrm{d}\hat{C}_t^0]}{\mathrm{d}t} = r(x_t - \hat{x}_t)^2$. In order to achieve the babbling cost, we need $x_t = \hat{x}_t$ almost surely, which means the optimal policy should be state-independent almost surely. Through (12), this requires $f - \gamma \star f$ to be a constant for almost all θ. When $\mu = 0$, this implies that $(\gamma \star f)'' = -\frac{2r}{\sigma^2}(f - \gamma \star f)$ is a constant almost everywhere. From Assumption 1, $\gamma \star f$ is twice differentiable, so that $(\gamma \star f)''$ is a constant, meaning $\gamma \star f(\theta) = \tilde{c}_0 + c_1\theta + c_2\theta^2$. This integral equation has the unique continuous solution $f(\theta) = \left(\tilde{c}_0 - \frac{c_2\sigma^2}{r}\right) + c_1\theta + c_2\theta^2$, where $\tilde{c}_0 - \frac{c_2\sigma^2}{r}$ can be denoted as c_0. Modification of the above on a zero-measure set generates an equivalence class.

For part (iii), the "if" direction can be verified by plugging $f(\theta) = c_0 + c_1\theta + c_2e^{-\frac{2\mu}{\sigma^2}\theta}$. The resulting action path $x_t = f(\theta_0) - c_1\mu t$ is deterministic. For the "only if" direction, repeat the same procedure in the proof of part (ii). When $\mu \neq 0$, this implies that $(\gamma \star f)'' + \frac{2\mu}{\sigma^2}(\gamma \star f)' = -\frac{2r}{\sigma^2}(f - \gamma \star f)$ is a constant, meaning $\gamma \star f(\theta) = \tilde{c}_0 + c_1\theta + c_2e^{-\frac{2\mu}{\sigma^2}\theta}$. This integral equation has the unique continuous solution $f(\theta) = \left(\tilde{c}_0 - \frac{c_1\mu}{r}\right) + c_1\theta + c_2e^{-\frac{2\mu}{\sigma^2}\theta}$, where $\tilde{c}_0 - \frac{c_1\mu}{r}$ can be denoted as c_0. Modification of the above on a zero-measure set generates an equivalence class.

Proof for Theorem 3

PROOF:

By definition, it suffices to show that

$$\frac{\partial x}{\partial \theta} = f'(\theta) - (\gamma \star f)'(\theta) = -\frac{2\mu(\gamma \star f)''(\theta) + \sigma^2(\gamma \star f)'''(\theta)}{2r}.$$

In equation (13), taking the derivative of $x_t(\theta^t)$ w.r.t. θ_t, one has

$$\left.\frac{\partial x_t}{\partial \theta_t}\right|_{\theta_t = \theta} = f'(\theta) - (\gamma \star f)'(\theta).$$

Also, by definition of γ, it is easy to verify that $g = \gamma \star f$ is a solution to the following ODE:

$$-\frac{\sigma^2}{2r}g'' - \frac{\mu}{r}g' + g = f. \tag{A6}$$

Therefore, further differentiation gives:

$$-\frac{\sigma^2}{2r}(\gamma \star f)'''(\theta) - \frac{\mu}{r}(\gamma \star f)''(\theta) = f'(\theta) - (\gamma \star f)'(\theta).$$

Proof for Proposition 2

PROOF:

To show (i), note that

$$\begin{aligned} \frac{\mathbb{E}_t\left[\mathrm{d}C_t^*\right]}{\mathrm{d}t} &= \frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}_t\left[(W_t - \gamma \star f(\theta_t))^2 + \frac{\sigma^2}{r}\gamma \star ((\gamma \star f)'^2)(\theta_t)\right] \\ &= \sigma^2 \gamma \star ((\gamma \star f)'^2)(\theta_t) \geqslant 0, \end{aligned}$$

where the second equality follows from Ito's lemma and the policy function, and the inequality comes from the fact that $(\gamma \star f)'^2 \geqslant 0$ and the convolution preserves the sign.

To show (ii), we start from the law of motion of W_t implied by the IC-FOC:

$$\mathrm{d}W_t = r(W_t - x_t)\mathrm{d}t = r(\gamma \star f(\theta_t) - f(\theta_t))\mathrm{d}t = \left(\frac{\sigma^2}{2}(\gamma \star f)''(\theta_t) + \mu(\gamma \star f)'(\theta_t)\right)\mathrm{d}t,$$

where the second equality holds by the policy function, and the third equality comes from (A6).

Relaxing the Strategy Set $\mathcal{L}$

PROOF:

The strategy set $\mathcal{L}$ limits the speed that the agent can lie in an exponential manner. It is assumed for technical simplicity. Now, I proceed to remove it. Without it, the global IC is problematic for some "crazy" strategies: lie exponentially at a very high rate. By doing that, the agent secures high flow payoffs at the cost of the continuation payoff that explodes to $-\infty$, *although* this does not happen on path. In the following, I construct a sequence of contracts that has a cost approaching C^*, so that the C^* in the main model is the infimum, not minimum.

Consider the optimal contract truncated at time T. Before the deadline T, do exactly as in the optimal contract. At time T, the action is frozen forever at $x_T = W_T$, so that the continuation payoff of the agent is promised even after the deadline. Obviously, with a finite deadline, the agent's infinite global scheme of deviation fails, since at the "Judgement Day" T, past deviations always factors in W_T which does not allow further Ponzi-like deviations.

I claim that this contract yields a cost C^T that approaches C^* as $T \to \infty$. At time T, the (positive) cost gap between the truncated contract and the optimal one is

$$\begin{aligned}
\Delta(\theta_T) &\equiv C^T(\theta_T, W_T) - C^*(\theta_T, W_T) \\
&= \gamma \star (W - f(\theta_T))^2 - (W - \gamma \star f(\theta_T))^2 - \frac{\sigma^2}{r}\gamma \star ((\gamma \star f)'^2)(\theta_T) \\
&= \gamma \star f^2(\theta_T) - (\gamma \star f(\theta_T))^2 - \frac{\sigma^2}{r}\gamma \star ((\gamma \star f)'^2)(\theta_T) \\
&\leqslant \gamma \star f^2(\theta_T) \\
&\leqslant 2\alpha_0^2 + \alpha_0^2 r \left(\frac{e^{-2\alpha_1\theta_T}}{r + 2\alpha_1(\mu - \alpha_1\sigma^2)} + \frac{e^{2\alpha_1\theta_T}}{r - 2\alpha_1(\mu + \alpha_1\sigma^2)} \right),
\end{aligned}$$

where the second inequality follows from Assumption 1. Hence,

$$\begin{aligned}
&C^T(\theta_0, W_0) - C^*(\theta_0, W_0) = e^{-rT}\mathbb{E}\left[\Delta(\theta_T)\right] \\
\leqslant\ & \alpha_0^2 e^{-rT} \left(\frac{e^{2\alpha_1(-\theta_0 - \mu T + \alpha_1\sigma^2 T)}}{r + 2\alpha_1(\mu - \alpha_1\sigma^2)} + \frac{e^{2\alpha_1(\theta_0 + \mu T + \alpha_1\sigma^2 T)}}{r - 2\alpha_1(\mu + \alpha_1\sigma^2)} \right),
\end{aligned}$$

where the last expression vanishes as $T \to \infty$, by Assumption 1.