

Dynamic Communication with Commitment

Yi Chen*
(*Job Market Paper*)

November 15, 2017

Abstract

I study the dynamic budgeting problem between the Headquarters and a division Manager, in which the Headquarters commits to a dynamic budgeting rule. The Manager privately knows the profitability of his own product when requesting budget from the Headquarters, but he always has a dominating personal motive to request more budget. With a theory model in continuous time, I first show that, unlike with one-shot interaction, a dynamic budgeting rule helps the Headquarters elicit and utilize information from the Manager in most cases. Second, in many situations the optimally allocated budget exhibits a *lagged* response to the change in the reported profitability. If the current reported profitability is high, the current budget will not increase, while the future budget will increase. Furthermore, the optimal rule may even stipulate *contrarian* budgeting: decrease the current budget when the reported profitability is high, and increase it when low. Finally, I show that the budget grows either monotonically tighter or looser over time. My results can be applied to more general principal-agent settings.

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

JEL Codes: D82, D83, D86

*Yale University, 28 Hillhouse Avenue, New Haven, 06511. Email: y.chen@yale.edu. I am grateful to my dissertation committee: Johannes Hörner, Larry Samuelson, and Dirk Bergemann. I also thank Yeon-Koo Che, Eduardo Faingold, Simone Galperti, Yuan Gao, Marina Halac, Nicolas Lambert, Elliot Lipnowski, Navin Kartik, Dmitry Orlov, Ennio Stacchetti, Philipp Strack, Zhe Wang, and Yiqing Xing for helpful comments and suggestions. All remaining errors are mine.

1 Introduction

Communication facilitates the transmission of information relevant for decision-making, the extent of which, however, is limited by a conflict of interests. The informed party (agent, he) often has a motive to misguide the uninformed party (principal, she), who takes actions based on the message conveyed by the former. In this paper, I investigate whether or not effective communication can be secured by a dynamic contract, and how to best elicit and utilize information from the agent.

As an example, consider the resource allocation office in the headquarters of a firm (principal), who must decide how to allocate resources over time to a division manager (agent) endowed with a unique product. From the perspective of the headquarters, the optimal amount of resources to be allocated to the division depends on the prospects of the product (profitability, technical parameters, etc.). However, having specialized in the product, the division manager understands its prospects much better. Hence, his knowledge is valuable to the headquarters. Communication problems arise if the manager's personal preferences are such that he always prefers more resources allocated to his division, *regardless* of the product's actual prospects. The headquarters is able to commit to a dynamic rule of resource allocation based on the manager's reports.¹ Considering the severe misalignment in preferences, is there any possibility for effective communication? If yes, what is the optimal rule of dynamic resource allocation? In particular, does the headquarters necessarily allocate more resources to the division when the manager reports better prospects? Does the agency problem pose an increasingly costly burden in the long run? Does the manager face tighter resource provisions over time?

State-independent preferences of the agent can arise in a number of situations. For example, a division manager holding empire-building motives always prefers more resources; a local government chasing political achievements prefers a larger infrastructure budget allocated by the central government. Given state-independent preferences (and a one-dimensional state), if the principal-agent relationship lasts for only *one period*, then there cannot be any effective communication. Indeed, as long as the contract specifies different actions upon different reports, the agent will always pick whichever report that induces the highest expected action. Hence, the expected action is not sensitive to the state, resulting in a severe loss of information.

The prospect for communication is greater when it involves *long-term* relationships with a changing state. In order to induce truthful reports from the agent, the principal no longer has to make the entire sequence of actions unresponsive to information. Instead, as long as

¹For example, a contract can be a total resource budget assigned to the manager across periods.

the *continuation* payoff of the agent is independent of the current report, incentives for truth-telling are provided. Since there are many different possible paths of actions that generate the same continuation payoff, the principal has degrees of freedom to be able to reallocate actions between the present and future in her favor. This is how a dynamic contract improves communication.

I discover a new channel, based on *information sensitivity*, that facilitates communication and shapes the contract in dynamic interactions. To explain this term, the headquarters-manager example is helpful. Define the headquarters' ideal amount of resource allocation as the *target*, which is a function of the current profitability. Then the slope of the target function, called the *information sensitivity*, captures how much the ideal amount of resources marginally changes with the profitability. The higher the information sensitivity, the more the headquarters' ideal resource allocation depends on the manager's private information.

Suppose the profitability of the product rises. In some circumstances, the ideal resource allocation increases considerably, but in other circumstances the increase is negligible. As an example of the latter case, if the prospect of the product is sufficiently good such that a marginal increase in the profitability does not justify any additional resources, then the information sensitivity is low. Similar situation also arises when the prospect of the product is bad enough, such that the ideal amount of resources cannot decrease further.

Information sensitivity determines the headquarters' trade-off between the present and the future, which in turn shapes the optimal contract. To optimally elicit and utilize the manager's private information about profitability, the headquarter carefully balances the current distortion in resource allocation with future distortions. If the information sensitivity is expected to be higher in the future than it is now, then future weighs more in the intertemporal trade-off. Since the total amount of resources must stay fixed to keep the incentives of the manager, the headquarter has to *decrease* the current resource allocation despite the increased profitability, so as to provide more resources in the future for a better match. The persistence in the state process is necessary to insure that higher current profitability augurs higher future profitabilities as well.

Formally, I solve for the optimal contract between a principal and an agent. The agent privately observes the evolution of a state process and continually reports to the principal, who in turn takes actions that affect the payoffs of both. The state evolves according to a Brownian motion. The agent can manipulate his report by inflating or shading the true process at any time. While the agent always prefers high actions regardless of the state, the principal's ideal action is state-dependent. Information is valuable for the principal, in that her flow cost is quadratic in the distance between the actual action and her ideal action. The latter, the *target*, is a function of the current state, the slope of which is the information

sensitivity mentioned before. The principal observes nothing except the agent's reports, and commits upfront to a dynamic contract specifying how actions are taken based on the report history. There are no monetary transfers.²

The model delivers two main results. The first pertains to the scope of communication: a dynamic contract enables effective communication if and only if the slope of the target function, or equivalently the information sensitivity, is nonlinear in the state. The second characterizes the optimal contract. Under certain conditions, the optimal rule stipulates *contrarian* budgeting: decrease the current resource provision when the reported profitability is high, and increase it when low. At the same time, the optimally allocated resources may exhibit a *lagged* response to the change in reported profitability. If the current profitability is claimed to be high, the current resources does not increase, but the *future* resources will. For the optimal contract to display those two properties, the necessary and sufficient condition is that the expected future information sensitivity is higher than its current counterpart, or equivalently, the expected future information sensitivity is convex in the current state. Hence, it is the sign of the third derivative, or *prudence*, of the target function that matters for the qualitative pattern of the optimal contract, and an analogy can be made with the literature on precautionary savings.

The intuition is as follows. A positive shock to the current state leads to higher future states in expectation, due to the persistence of the state. Ideally the principal would want to increase actions both now and in the future, but this violates the incentive constraints of the agent. In order to preserve incentives, the agent's total payoff (the principal's total action) should be unresponsive to the current report. Hence, a higher current action must be accompanied by lower future actions, and vice versa. The trade-off depends on the slope of the target function, i.e., the aforementioned information sensitivity. If the expected future information sensitivity is convex, then due to the random evolution of the state, the future information sensitivity is higher in expectation than it is now. In other words, the expected future target increases *relative* to the current target. Hence, with a positive state shock, the principal optimally lowers the current action in exchange for higher actions in the future to rebalance distortions across periods, even though it entails a contrarian action that works against her preferences. In this way, while taking actions that seemingly move against the target, the principal's actions are *correct on average*. The gains from this on-average correctness more than compensate the losses from the contrarian nature of the action. The trade-off is reversed if the slope of the expected future target is concave, in which case the action moves in the same direction as the target. When the slope is linear in the state,

²I show in Section 6.2 that when there is limited liability for the agent, allowing for monetary transfers do not affect the qualitative results.

future is exactly as important as the present, and hence there is no profitable trade-off in any direction. In this case, communication fails, although unlike one-shot interactions, the principal *can* make the contract responsive to reports, but she chooses not to due to her preferences.

The contrarian property of the contract relies on the persistence of the state process. If the state is i.i.d. over time, then the distribution of future states does not depend on the current state, and there is no reason to sacrifice current action in exchange for a better match in the future. This also sheds light on the reason why contrarian contract does not arise in the literature where the state process is modeled as a two-state Markov chain. On one hand, with only two states, the target function is linear by definition. On the other, mean-reversion is automatically built in with the Markov chain, so that the persistence of the state is weaker. Combining these two factors, the future is always less important than the present, and hence the action always moves along with the target.

Finally, the model delivers predictions regarding the long-term behavior of the contract and payoffs. The principal's continuation cost follows a sub-martingale, as the distortion from the incentive constraint accumulates over time. The agent's continuation payoff drifts monotonically to infinity if the target is convex, and to minus infinity if the target is concave. Again, it is the shape of the target function that determines the evolution of the principal's payoff.

This paper is connected to the literature of communication. Since Crawford and Sobel (1982) and Green and Stokey (2007), there is a large body of literature on cheap talk with a fixed one-dimensional state (Aumann and Hart (2003), Krishna and Morgan (2001, 2004), Goltsman, Hörner, Pavlov, and Squintani (2009), etc.). In these papers the static nature of decision requires significant congruence of preferences in order for informative equilibria to exist. Since commitment power is lacking in cheap talk, the commitment by contract in my paper brings it closer to the literature of delegation (Holmström (1977), Alonso and Matouschek (2008), Amador and Bagwell (2013)) studies communication problems where the principal commits to an action set and the agent takes whichever action he likes within this set.

More closely related is the literature on multi-dimensional cheap talk and dynamic allocation problems. Battaglini (2002) and Chakraborty and Harbaugh (2010) explore the possibility of one-shot communication with higher-dimensional states, and find qualitatively different patterns of communication than in the one-dimensional case. They show that equilibria with meaningful communication generically exist. Golosov, Skreta, Tsyvinski, and Wilson (2014) extend the cheap talk game of Crawford and Sobel (1982) to multiple periods with a fixed state, and find that communication may improve in later periods. Jackson

and Sonnenschein (2007) consider the problem of how to link independent replicas. They introduce a quota mechanism that takes advantage of the Law of Large Numbers to achieve asymptotic efficiency. Antič and Steverson (2016) feature a static mechanism with multiple agents, each having state-independent preferences over his own allocation. They find that the optimal mechanism may display “strategic favoritism” to exploit the super-modularity of productivity among agents. Renault, Solan, and Vieille (2013) study a repeated cheap talk game where the state follows a finite-state Markov chain. They construct quota-like equilibria to establish the limit set of payoffs when players become infinitely patient. Guo and Hörner (2017) investigate the optimal allocation mechanism without transfer, where there are binary actions and binary persistent states. They describe the optimal mechanism in terms of a generalized quota, and find asymptotics different from immiseration. In these papers, counter-cyclical action does not arise for various reasons: the state is fixed, the states are independent, or the state is binary. My paper features an imperfectly persistent state where the information sensitivity varies from state to state. As a result, I find conditions where it is optimal to *save* quota when it is otherwise tempting to use it. Moreover, the evolution of payoffs is driven by a new force, i.e., the differential information sensitivity of the principal.

This paper is also related to the literature on dynamic agency problems with transfer. For instance, Sannikov (2008) studies a dynamic moral hazard problem without private information, and Demarzo and Sannikov (2016) and He, Wei, Yu, and Gao (2017) focus on the dynamic interactions between hidden action and private learning. The role of persistent private information in a dynamic taxation/subsidy mechanism is extensively explored in Fernandes and Phelan (2000), Williams (2011) and Kapička (2013). The absence of transfers in my model generates quite different implications for the optimal contract, although, as is shown in the extension, transfers with limited liability partially preserve the results from the main model.

The remainder of the paper is organized as follows. Section 2 presents a two-period example to illustrate the key trade-off in the optimal contract. Section 3 lays out the setting for the continuous-time model. Section 4 simplifies the problem through the revelation principle and derives a necessary condition and a sufficient condition for incentive compatibility. Section 5 fully analyzes the optimal contract and gives implications and applications. Section 6 discusses some extensions, and Section 7 concludes.

2 A Two-Period Example

To illustrate the essential role of inter-temporal trade-offs that shape a dynamic contract, it is convenient to start with a two-period example. There are two periods: $t = \{1, 2\}$. A state $(\theta_t)_{t=1,2}$ follows a 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 $t = 1, 2$, 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 . There is no transfer. The principal's cost in period t is $(x_t - f(\theta_t))^2$, and the agent's payoff is simply x_t . The function $f(\cdot)$ is called the *target* function throughout the entire paper, characterizing the ideal action at every state. The total cost or payoff is the sum over both periods, with a discount factor $\delta > 0$ on the second period.

A dynamic contract is a pair $(x_1(\hat{\theta}_1), x_2(\hat{\theta}_1, \hat{\theta}_2))$, mapping report histories into actions. I focus on contracts that induce on-path truth-telling for the agent. The principal's problem is written as:

$$\begin{aligned} \min_{x_1(\cdot), x_2(\cdot, \cdot)} \quad & \mathbb{E} [(x_1 - f(\theta_1))^2 + \delta(x_2 - f(\theta_2))^2] \\ \text{s.t.} \quad & x_1(\theta_1) + \delta \mathbb{E} [x_2(\theta_1, \theta_2) | \theta_1] \geq x_1(\hat{\theta}_1) + \delta \mathbb{E} [x_2(\hat{\theta}_1, \hat{\theta}_2) | \theta_1] \quad \forall \theta_1, \hat{\theta}_1, \hat{\theta}_2, \quad (1) \\ & x_2(\theta_1, \theta_2) \geq x_2(\theta_1, \hat{\theta}_2) \quad \forall \theta_1, \theta_2, \hat{\theta}_2. \quad (2) \end{aligned}$$

Constraint (2) requires that truth-telling is optimal for the agent in period 2 after a truthful report in period 1. Since the agent's payoff is completely state-independent, (2) implies $x_2(\theta_1, \theta_2) \geq x_2(\theta_1, \hat{\theta}_2) \geq x_2(\theta_1, \theta_2)$ for any θ_2 and $\hat{\theta}_2$. As a result, $x_2(\theta_1, \theta_2)$ must be independent of θ_2 , and I abuse notation by writing $x_2(\theta_1)$ for short.

Constraint (1) governs the period-1 incentive, and states that the agent obtains the highest expected total payoff by truth-telling in both periods, among all reporting strategies. Since x_2 does not depend on θ_2 , (1) is simplified to $x_1(\theta_1) + \delta x_2(\theta_1) \geq x_1(\hat{\theta}_1) + \delta x_2(\hat{\theta}_1)$ for all $\theta_1, \hat{\theta}_1$. Again, due to the state independence, this must hold with equality for any pair of states, and (1) finally reduces to:

$$x_1(\theta_1) + \delta x_2(\theta_1) = \text{constant}.$$

The constant on the right-hand side is obviously interpreted as the total payoff of the agent,

and is hence denoted as W . The optimal level of W is endogenously chosen by the principal as part of the maximization problem.

Before solving the optimal contract, it is important to briefly discuss the role of dynamics. If the interaction lasts for only one period, there is no hope of meaningful communication. To see why, notice that incentive compatibility requires the contract to specify an action independent of the reported state, otherwise the agent would always select the report that induces the highest action. Therefore, information is wasted and communication fails because the constant action can be taken without information at all. The logic changes when there are multiple periods. The period-1 IC (1) only requires the total payoff $x_1 + \delta x_2$ of the agent to be independent of $\hat{\theta}_1$, but the pair of actions (x_1, x_2) still has one degree of freedom to adjust. While the agent is indifferent about these adjustments, the principal, who has different preferences, values the ability to reallocate actions between periods in response to $\hat{\theta}_1$. As a result, the optimal contract may perform better than babbling, i.e., taking an action sequence independent of any information.

Replacing $x_2(\theta_1)$ by $\frac{W-x_1(\theta_1)}{\delta}$ and plugging it into the original problem, one arrives at an unconstrained version with $x_1(\cdot)$ and W as variables. The optimal contract is obtained by first finding the optimal action sequence given any constant W , and then optimizing over W (algebra relegated to Appendix). The solution reads (with $\hat{\theta}_1 = \theta_1$):

$$x_1^*(\theta_1) \equiv \frac{\delta}{1+\delta}(f(\theta_1) - \mathbb{E}[f(\theta_2)|\theta_1]) + \frac{\mathbb{E}f(\theta_1) + \delta\mathbb{E}f(\theta_2)}{1+\delta}, \quad (3)$$

$$x_2^*(\theta_1) \equiv \underbrace{-\frac{1}{1+\delta}(f(\theta_1) - \mathbb{E}[f(\theta_2)|\theta_1])}_{\text{responsive to } \theta_1} + \underbrace{\frac{\mathbb{E}f(\theta_1) + \delta\mathbb{E}f(\theta_2)}{1+\delta}}_{\text{independent of } \theta_1}. \quad (4)$$

Note that the second term of x_1^* or x_2^* is the same constant. Actually, $W^* \equiv \mathbb{E}f(\theta_1) + \delta\mathbb{E}f(\theta_2)$ is the optimal choice for the total payoff W . Since $x_1^* + \delta x_2^* = W^*$, the total payoff is shared evenly across periods as annuities.

More importantly, the first terms in x_1^* and x_2^* represent the response of the action to the first report. The period-1 incentive constraint $x_1 + \delta x_2 = W$ acts like a “budget” or “quota” for the inter-temporal allocation of actions. Within this budget set, x_1^* and x_2^* optimally reacts to information for cost minimization purposes. Specifically, differentiating both sides of (3) with respect to θ_1 , we have:

$$\frac{dx_1^*}{d\theta_1} = \frac{\delta}{1+\delta} \left(f'(\theta_1) - \frac{d}{d\theta_1} \mathbb{E}[f(\theta_2)|\theta_1] \right). \quad (5)$$

The sign of this derivative is not entirely obvious. Suppose f is increasing, so that the target




Figure 1: Iso-cost curves and the budget line. Left panel: the family of iso-cost curves and the budget line. Middle panel: increase in $f(\theta_1)$ causes a decrease in x_1 . Right panel: increase in $f(\theta_1)$ causes an increase in x_1 .

$f(\theta_1)$ goes up with the state θ_1 . However, due to the “budget constraint,” an increase in x_1 must be accompanied by a decrease in x_2 , causing x_2 to move in the opposite direction of the change in the expected period-2 target $\mathbb{E}[f(\theta_2)|\theta_1]$. This effect is reflected in the second term in the bracket of (5).

If the sign of $f'(\theta_1)$ is reversed by subtracting $\frac{d}{d\theta_1}\mathbb{E}[f(\theta_2)|\theta_1]$ from it, then the action x_1^* and the target $f(\theta_1)$ move in opposite directions, which I call *counter-cyclical* action. If the sign is not reversed so that the action and the target move together, then the action is called *pro-cyclical*. For example, when $f(\theta) = e^\theta$, the action is counter-cyclical for all θ_1 . For the target function $f(\theta) = \theta - \frac{1}{2}|\theta|$, counter-cyclical action occurs only at a subset of states. The action is always pro-cyclical if $f(\theta) = \frac{|\theta|^{5/2}}{\theta}$. 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 sequence does not respond to information at all. Since information-independent actions can always be taken without the presence of the agent, communication is considered to be a failure.

Figure 1 gives a geometric illustration of the cost minimization problem. Incentive compatibility pins the pair (x_1, x_2) on a budget line, as shown in the left panel. The concentric ellipses are the principal’s iso-cost curves. The state θ_1 determines the center of the ellipses, with coordinates given by the current target $f(\theta_1)$ and the conditional expectation of the target in period 2, $\mathbb{E}[f(\theta_2)|\theta_1]$. 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. The other two panels show how the tangent point moves with θ_1 . Suppose f is increasing so that a higher θ_1 leads to a higher $f(\theta_1)$ and, due to the persistence of the state, a higher expected target $\mathbb{E}[f(\theta_2)|\theta_1]$ as well. Graphically, a positive shock in θ_1 causes the center of ellipses to move from O_A

to O_B . If the point O_B lies above the line O_AA , then the new tangent point B is to the northwest of the old tangent point A , resulting in a decline in x_1 (see the middle panel of Figure 1). Hence, when $\frac{d\mathbb{E}[f(\theta_2)|\theta_1]}{df(\theta_1)} > 1$, i.e., the expected future target is more sensitive than the current target to a state shock, counter-cyclical action occurs. Otherwise, the action is pro-cyclical (see the right panel). This condition is verified in (5).

Some comments follow. First, persistence in the state process is necessary for counter-cyclical actions to appear. Without persistence, O_B moves horizontally to the right of O_A , lying below O_AA . Second, the optimal contract features a “quota” mechanism, but instead of draining the quota when it is most tempting, the optimal contract may do the opposite: that is, save the quota (delay its use) when the state increases and use it when the state decreases. Finally, the period-2 information cannot be used in this example because the second period is the last; as the horizon increases, there are more periods in which information is utilized.

In the next section, I present the settings for the full model with continuous time and infinite horizon. Continuous time allows for the smooth evolution of information and convenient analysis of conditions for actions to be pro/counter cyclical. It also enables the study of the asymptotics of the contract.

3 Continuous-Time Model: Settings

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

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

where $\mathbf{Z} = (Z_t)_{t \geq 0}$ is the standard Brownian motion on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$, and $\sigma > 0$ is a constant.³ The bold letter $\boldsymbol{\theta}$ represents the state process while a plain letter θ stands for a generic value of the state (the same notation applies to all other processes). The initial state θ_0 is common knowledge.

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

$$d\hat{\theta}_t = l_t dt + d\theta_t,$$

where l_t , chosen by the agent at every moment $t \geq 0$, is interpreted as the “speed of lying.” In other words, the agent can alter the true state process by adding a drift of his choice. The principal takes action x_t at every $t \geq 0$.

³For simplicity, the state process does not have a drift. The case of constant drift or mean-reversion is examined in Section 6.1.

Interests are misaligned in the following sense. While the principal’s favorite action depends on the state, the agent only wishes to induce actions to be as high as possible. Specifically, the principal’s flow cost from a state-action pair (θ, x) is $(x - f(\theta))^2$. It takes the form of a quadratic cost resulting from the gap between the actual action x_t and the *target* action $f(\theta_t)$. This quadratic cost structure can be justified as a reduced-form payoff of the principal: action x generates a linear benefit $f(\theta)x$ but entails a quadratic cost $\frac{1}{2}x^2$. The net flow profit is then maximized at $x^* = f(\theta)$, and any action x other than x^* yields a relative cost:

$$\left(f(\theta)x^* - \frac{1}{2}x^{*2}\right) - \left(f(\theta)x - \frac{1}{2}x^2\right) = \frac{1}{2}(x - f(\theta))^2,$$

which is exactly the quadratic cost assumed above with some re-scaling. In general cases, the quadratic cost is often a good approximation.

The agent’s flow payoff is simply x_t , independent of the state.⁴ The risk neutrality in the agent’s payoff is assumed for convenience. On the one hand, the payoff from this interaction may consist of only a small fraction in the agent’s utility, and risk aversion hardly displays. On the other hand, even if the agent’s payoff is an increasing but non-linear function of the action, one can arguably re-normalize the action to be the payoff itself.⁵ The state-independent payoff renders communication babbling (i.e., where information does not influence actions) in a one-shot interaction, as the agent always reports the state that induces the highest action.

The players share the same discount rate $r > 0$. In the following, I write $\theta^t \equiv (\theta_s)_{0 \leq s \leq t}$ as the state history up to time t , and when $t = \infty$, it represents the entire path of state (a similar convention is used for the report history and action history). Fixing x^∞ and θ^∞ , the paths of action and state, the realized total cost of the principal and total payoff of the agent are, respectively:

$$u_R(x^\infty, \theta^\infty) = \int_0^\infty r e^{-rt} (x_t - f(\theta_t))^2 dt,$$

$$u_S(x^\infty, \theta^\infty) = \int_0^\infty r e^{-rt} x_t dt,$$

whenever well-defined.

The principal does not observe the state process except for its initial value θ_0 . Further-

⁴This insatiable preference of the agent excludes the possibility of “moderate” bias as in Crawford and Sobel (1982), although the main results do not rely on this extreme specification. Starting with flow payoff $-\frac{1}{2b}(x_t - f(\theta_t) - b)^2 + \frac{b}{2}$ for the agent and letting $b \rightarrow \infty$, we have the limiting flow payoff $x_t - f(\theta_t)$. The term $-f(\theta_t)$ is exogenous and can be left out.

⁵With this re-normalization the principal’s cost is no longer exactly quadratic, but still good enough as an approximation.

more, I assume that the principal does not observe her own flow payoffs or any signals of past states, precluding the possibility of inference. By contrast, the agent observes the state path as it evolves.

The principal commits to a contract at time zero. A contract is a $\hat{\theta}$ -measurable process \mathbf{x} specifying an action $x_t(\hat{\theta}^t) \in \mathbb{R}$ as a function of the report history, for all $t \geq 0$. There is no transfer of money. A strategy of the agent is a θ -measurable process \mathbf{l} . It prescribes the lying behavior $l_t(\theta^t)$ of the agent as a function of the state history for all $t \geq 0$. I define the space of feasible strategies as:

$$\mathcal{L} \equiv \left\{ l : \lim_{t \rightarrow \infty} e^{-2rt} \mathbb{E} \left[e^{\frac{1}{2\sigma^2} \int_0^t l_s^2 ds} \right] = 0 \right\},$$

to exclude explosive strategies in the limit. This assumption is for simplicity only, but is not essential.⁶ The agent has no outside option.⁷

Given a contract-strategy pair (\mathbf{x}, \mathbf{l}) , the total expected cost and payoff are respectively:

$$U_R(\mathbf{x}, \mathbf{l}) = \mathbb{E}^l \left[\int_0^\infty r e^{-rt} (x_t - f(\theta_t))^2 dt \right],$$

$$U_S(\mathbf{x}, \mathbf{l}) = \mathbb{E}^l \left[\int_0^\infty r e^{-rt} x_t dt \right],$$

whenever well-defined, where \mathbb{E}^l denotes the expectation under strategy \mathbf{l} . Hereafter, “pay-off” 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 \mathbf{l} to maximize his payoff given contract \mathbf{x} . The principal designs a contract \mathbf{x} to minimize her cost given the agent’s optimal choice of strategy in reaction to the contract. In the case where the agent has multiple optimal strategies given a contract, I assume the tie is always broken in the principal’s favor.

4 Incentives of the Agent

This section describes two simplifying procedures leading to the derivation of the optimal contract in Section 5. First, from amongst all contracts, I restrict attention to those that implement truth-telling from the agent by presenting a version of the Revelation principle. Second, I use the first-order approach to derive a necessary condition for incentive compatibility, and then give a sufficient condition. In Section 5, the necessary condition is used in

⁶In the end of Appendix, I show the effect of relaxing the strategy set.

⁷An ex ante participation constraint can be easily introduced. See the extension in Section 6.3.

place of the incentive constraints, yielding a candidate solution to the *relaxed* problem. The solution is then verified to satisfy the sufficient condition.

4.1 Revelation Principle

A strategy is called *truthful* if it is identically zero, denoted as \mathbf{l}^\dagger . A contract \mathbf{x} is *truthful* if the agent maximizes his payoff by taking the truthful strategy. By Lemma 1 below, I focus on truthful contracts without loss of generality.

Lemma 1 (Revelation Principle)

For any contract \mathbf{x} that induces agent's strategy \mathbf{l} , there exists a truthful contract \mathbf{x}^\dagger that implements the same mapping from state paths θ^∞ into action paths x^∞ .

Proof. See Appendix. ■

Among truthful contracts (“truthful” will be omitted henceforth), the principal’s problem is written as:

$$\min_{(x_t)_{t \geq 0}} \mathbb{E}^{\mathbf{l}^\dagger} \left[\int_0^\infty r e^{-rt} (x_t - f(\theta_t))^2 dt \right] \quad (6)$$

$$\text{s.t.} \quad \mathbb{E}^{\mathbf{l}^\dagger} \left[\int_0^\infty r e^{-rt} x_t dt \right] \geq \mathbb{E}^{\mathbf{l}} \left[\int_0^\infty r e^{-rt} x_t dt \right], \quad \forall \mathbf{l} \in \mathcal{L}. \quad (7)$$

The incentive constraints (7) guarantee that any strategy \mathbf{l} achieves, at most, the payoff from truth-telling \mathbf{l}^\dagger . The constraints are expressed as of time zero. However, they also guarantee incentive compatibility at all later $t > 0$ since the agent faces a decision problem with time-consistent preferences. Indeed, if a strategy’s continuation is suboptimal on a set of histories with positive probability, the strategy itself is suboptimal at time zero.

Hidden behind (7) is the presumption that the contract is such that the payoff of the agent is well-defined for any \mathbf{l} . As explained below, this is without loss of generality.

4.2 Incentive Compatibility: Necessary Condition

The first-order approach in the literature (Williams (2011), Kapička (2013), Demarzo and Sannikov (2016), etc.) derives a local version of the incentive constraints, namely, conditions under which the agent does not profit by locally deviating from truth-telling.

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

$$W_t(\mathbf{x}) \equiv \mathbb{E} \left[\int_t^\infty r e^{-r(s-t)} x_s ds \middle| \mathcal{F}_t \right],$$

as the agent’s *on-path* expected continuation payoff conditional on the filtration \mathcal{F}_t , evaluated at time t . As verified in Section 5 after solving the optimal contract, this is the only state variable that summarizes the public history. Henceforth, I suppress the dependence of W_t on \mathbf{x} to simplify notations.

The use of the continuation payoff as one of the sufficient statistics for the entire history is common when it comes to equilibrium payoffs (Abreu, Pearce, and Stacchetti (1986), Thomas and Worrall (1990), etc.), but in a setting where the agent has *persistent* private information, an additional variable is often needed. For example, Williams (2011) and Kapička (2013) use the *marginal* continuation payoff as an extra state variable in order to govern the agent’s incentive gap caused by previous deviations. Guo and Hörner (2017) select the vector of conditional continuation payoffs as an alternative set of two sufficient statistics, which can be considered as a discrete version of the continuation payoff and the marginal continuation payoff. In my model, the marginal continuation payoff ceases to be a state variable even if the private information is persistent, because the agent’s payoff is independent of the state and the state evolution has independent increments. To see why, notice that 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. In the proof of Proposition 1, it is formally shown that the marginal continuation payoff is a constant zero.

Given a contract \mathbf{x} , the *on path* evolution of \mathbf{W} can be written as a diffusion according to Lemma 2.

Lemma 2 (Martingale Representation Theorem)

Given a contract \mathbf{x} , there exists a $\hat{\boldsymbol{\theta}}$ -measurable process $\boldsymbol{\beta} = (\beta_t)_{t \geq 0}$ such that:

$$dW_t = r(W_t - x_t)dt + r\beta_t(\underbrace{d\hat{\theta}_t}_{=\sigma dZ_t}). \tag{8}$$

Proof. See Appendix. ■

The multiplier $r\beta_t$ in (8) is interpreted as the instantaneous slope of the continuation payoff with respect to reported states, or “strength of incentives” (He, Wei, Yu, and Gao (2017), Demarzo and Sannikov (2016)). On path, $\hat{\boldsymbol{\theta}} = \boldsymbol{\theta}$, so that the second term on the right-hand side has a zero mean. The first term represents the drift of W_t : it grows at interest rate r and drops as flow payoff x_t is paid out to the agent, consistent with the way W_t is defined.

The second term, the diffusion, governs the incentives. Given an instantaneous slope $r\beta_t$,

the agent can manipulate his continuation payoff by misreporting.⁸ Suppose $\beta_t > 0$ and, unbeknownst to the principal, the agent lied and increased $d\hat{\theta}_t$ above its true value, then the principal would simply interpret this as a large realized increment in θ_t , and consequently promise the agent a higher continuation payoff W_t according to the contract. If $\beta_t < 0$, then lying in the opposite direction becomes profitable. Therefore, in order to deter a local deviation from truth-telling, this slope $r\beta_t$ must be identically zero. Proposition 1 shows that the coefficient β being identically zero is indeed a necessary condition for incentive compatibility.

Proposition 1 (IC-Necessity)

A necessary condition for incentive compatibility is $\beta_t = 0$ for all $t \geq 0$, a.s. \mathbb{P} .

Proof. See Appendix. ■

The proposition presents a simple and intuitive necessary condition for incentive compatibility: the only way to induce truth-telling from a state-independent agent is to entirely disentangle his future payoffs from his reports. The promised future payoffs are thus unresponsive to current information shocks, but the payoff process is nonetheless stochastic because its drift depends on current actions, which can be stochastic (see (8)).

As foreshadowed in the two-period model, this necessary condition reflects the limitation of what the principal can do in a dynamic setting. An interesting comparison can be made with a static setting, in which there is only a fixed one-dimensional state and a one-shot interaction between the two parties. In this case, the state-independent preferences of the agent destroy any scope of communication because the contract must assign a constant action to all possible states in order to elicit truth, which in turn implies the complete dissipation of information value. In my dynamic model, the limitation caused by incentive compatibility is still severe, in that the *continuation payoff* remains unresponsive to information. However, some degree of freedom exists since a fixed expected continuation payoff can be supported by different paths of actions. The choice of different action paths within the degree of freedom reflects the partially effective usage of information. In Section 5, I formally elaborate the optimal utilization of information through a dynamic contract.

4.3 Incentive Compatibility: Sufficient Condition

The state-independent payoff of the agent provides a simple structure of off-equilibrium payoffs, which leads to the following sufficient condition for incentive compatibility.

⁸In this continuous-time model, the continuation payoff is the objective for maximization for the agent because the weight of flow payoff is literally zero.

Proposition 2 (IC-Sufficiency)

If $\beta_t = 0$ for all $t \geq 0$ a.s. \mathbb{P} , and there exists constants $c_0 > 0, c_1 \in (0, \sqrt{\frac{r}{2\sigma^2}})$ such that $|x_t - W_t| \leq c_0(e^{c_1\theta_t} + e^{-c_1\theta_t})$ for all t , then the incentive constraints (7) are satisfied.

Proof. See Appendix. ■

From Proposition 2, the only additional requirement from necessity to sufficiency is the condition on the growth rate of $x_t - W_t$ with respect to the state. As long as this term does not grow too explosively, the on-path and off-path continuation payoffs satisfy the transversality condition. Actually, the total payoff is a constant for any strategy $\mathbf{l} \in \mathcal{L}$. The state-independent preferences of the agent contribute to this desirable feature, because the agent's actual payoffs are observable despite his ability to lie about the state.

5 Optimal Contract

Before solving the optimal contract, I first analyze the complete information benchmark and present the optimal contract without respecting the incentive constraints. Then I return to the private information case, find the optimal contract, and derive its properties.

A weaker version of the incentive constraints has been given by Proposition 1. The principal is said to solve the relaxed problem if she minimizes her cost subject to this necessary condition. The candidate solution is then verified to satisfy the sufficient condition in Proposition 2.

From this section on, I pose some regularity conditions on the target function $f(\cdot)$:

Assumption 1 (Regularity)

- (i) The target $f(\cdot)$ is piecewise \mathcal{C}^2 ;
- (ii) There exists $\alpha_0 > 0, \alpha_1 \in [0, \sqrt{\frac{r}{2\sigma^2}})$ such that:

$$|f(\theta)| \leq \alpha_0(e^{\alpha_1\theta} + e^{-\alpha_1\theta}).$$

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, which is necessary for costs and payoffs to be finite.

5.1 Optimal Decision with Complete Information

In the complete information case, the principal observes the state process herself. The incentives of the agent are completely ignored, but the optimal policy is derived for any

payoff promises that the principal plans to deliver to the agent.

Consider the principal's problem, subject to the agent's payoff being W_0 :

$$\begin{aligned} \underline{C}(\theta_0, W_0) &\equiv \min_x \mathbb{E} \left[\int_0^\infty r e^{-rt} (x_t - f(\theta_t))^2 dt \right] \\ \text{s.t. } W_0 &= \mathbb{E} \left[\int_0^\infty r e^{-rt} x_t dt \right]. \end{aligned} \quad (9)$$

With some algebra (relegated to Appendix), the solution reads:

$$\underline{x}_t(\theta_t, W_0) = f(\theta_t) + \left(W_0 - \mathbb{E} \left[\int_0^\infty r e^{-rs} f(\theta_s) ds \right] \right), \quad (10)$$

$$\underline{C}(\theta_0, W_0) = \left(W_0 - \mathbb{E} \left[\int_0^\infty r e^{-rs} f(\theta_s) ds \right] \right)^2. \quad (11)$$

The solution is intuitive. The optimal policy has two parts. The first part $f(\theta_t)$ simply serves to track the target action one-to-one. The second part (with brackets) is a *constant* adjustment to bring the agent's payoff to the promised level: if the principal were to take actions to match the target, then the agent would obtain a payoff $\mathbb{E} \left[\int_0^\infty r e^{-rs} f(\theta_s) ds \right]$ in expectation. In order to make up for the difference between this and the promised payoff W_0 , the principal, who has a quadratic loss function, optimally chooses a constant adjustment to smooth the distortions.

The term $\mathbb{E} \left[\int_0^\infty r e^{-rs} f(\theta_s) ds \right]$ can be further simplified. Using Fubini Theorem, the order of integrals can be switched, and we have:

$$\mathbb{E} \left[\int_0^\infty r e^{-rs} f(\theta_s) ds \right] = \gamma \star f(\theta_0),$$

where γ is the Laplace density function with parameter $\left(0, \frac{\sigma}{\sqrt{2r}}\right)$:

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

and $\gamma \star f$ is the convolution $\gamma \star f(\theta) \equiv \int_{-\infty}^\infty \gamma(z) f(\theta - z) dz$. Hence, the economic interpretation of the convolution is simply the *expected discounted future target*. See Figure 2 for the shape of γ and the illustration of the convolution imposed on a sample target function. The convolution is extensively used below.

The minimal cost \underline{C} is a quadratic function in W_0 . This quadratic form is inherited from the flow cost along with the fact that the burden of promising a sub-optimal level of




Figure 2: Left panel: the Laplace density function γ . Right panel: original target function f (solid curve) and its transform $\gamma \star f$ (dashed curve).

payoff is smoothed over time by keeping a constant distance from the target. Moreover, if the principal is free to choose the initial promise W_0 , then the cost is further reduced to 0 at $W_0 = \gamma \star f(\theta_0)$. In that case, the principal simply matches the action x_t with the target $f(\theta_t)$ at all times, and by doing so she automatically guarantees a payoff $\gamma \star f(\theta_0)$ to the agent. Therefore, one can claim that the cost \underline{C} is entirely from *promise keeping*—fulfilling a promised payoff that is too high or too low. When the state is the agent’s private information, the principal can no longer perfectly match the action with the state. Instead, she must trade-off between the match of current action versus the match of promised payoff in the future, a key argument in the next subsection.

5.2 Optimal Contract with Private Information

Here I proceed to using a recursive formulation to solve the relaxed problem, with only the necessary condition of incentives derived in Section 4.2. Conjecture that the optimal contract can be written in terms of the two state variables: the state θ , which summarizes the agent’s private history by Markov property of Brownian motion; and the continuation payoff W , which summarizes the public history. In Appendix, I formally prove that the candidate contract derived from this conjecture is indeed the solution to the original problem (6).

If θ and W are, as conjectured, sufficient for the continuation contract, then the lowest cost achievable by the principal, $C(\theta, W)$, also has the two variables as arguments. The evolution of the on-path continuation payoff process \mathbf{W} is derived in (8), and I copy it here:

$$dW_t = r(W_t - x_t)dt + r\beta_t\sigma dZ_t = r(W_t - x_t)dt,$$

where the second equality follows from the necessary condition $\beta_t = 0$.

With the law of motion in θ and W , the cost function must satisfy the following functional equation:

$$rC(\theta, W) = \min_x r(x - f(\theta))^2 + r(W - x)C_W(\theta, W) + \frac{\sigma^2}{2}C_{\theta\theta}(\theta, W). \quad (12)$$

The right-hand side of the equation consists of three terms: the normalized flow cost; the expected change in the future cost due to the drift in W ; and the cost change due to the volatility in θ . Other than (12), the cost function must also satisfy the transversality condition:

$$\lim_{t \rightarrow \infty} e^{-rt} \mathbb{E}[C(\theta_t, W_t)] = 0, \quad (13)$$

so that the cost C is consistent with the action path. Theorem 1 selects the unique solution to (12) and (13), and verifies that it achieves the minimum cost in the original problem (6).

Theorem 1 (Cost Function)

By promising the agent a payoff of W_0 , the principal's minimum cost is:

$$C^*(\theta_0, W_0) \equiv (W_0 - \gamma \star f(\theta_0))^2 + \frac{\sigma^2}{r} \gamma \star (\gamma \star f')^2(\theta_0). \quad (14)$$

Proof. See Appendix. ■

The proof of the theorem takes two steps. In step one, the martingale verification is performed to show that the cost function C^* is the lowest cost obtainable in the relaxed problem. In step two, I proceed to verify that there exists a contract achieving C^* which satisfies the sufficient condition in Proposition 2. Hence, C^* is the solution to (6).

Theorem 1 has two immediate implications with regard to the cost function. First, for any initial state θ_0 , the cost function is minimized at $W_0 = \gamma \star f(\theta_0)$, exactly the same as with complete information. This implies that the incentive problem serves not to distort the actions *on average*, but to affect the way in which the action responds to the state. Second, by comparing C^* to its complete-information counterpart \underline{C} , one finds the difference to be the second term on the right-hand side of (14), which is non-negative. Clearly, this extra cost is interpreted as the *incentive provision* cost due to private information.

The policy function is obtained from the cost function, with θ and W as arguments. In order to eventually express the contract as a function of the report history only, I insert the evolution of \mathbf{W} into the policy function to eliminate its dependence on W .

Proposition 3 (Optimal Contract)

The policy function reads:

$$x^*(\theta, W) \equiv W + f(\theta) - \gamma \star f(\theta). \quad (15)$$

The optimal contract \mathbf{x}^* specifies:

$$x_t^*(\theta^t) \equiv f(\theta_t) - \gamma \star f(\theta_t) + \gamma \star f(\theta_0) + \frac{\sigma^2}{2} \int_0^t (\gamma \star f)''(\theta_s) ds; \quad (16)$$

if f'' always exists, then equivalently

$$x_t^*(\theta^t) = f(\theta_0) + \int_0^t \frac{\sigma^2}{2} f''(\theta_s) ds + \int_0^t [f'(\theta_s) - (\gamma \star f)'(\theta_s)] d\theta_s. \quad (17)$$

Proof. See Appendix. ■

The policy function (15) has three components: the “annuity” from the continuation payoff W , the “first-best” response to information, $f(\theta)$, and the adjustment due to incentive constraints, $-\gamma \star f(\theta)$. The first component guarantees that the policy respects the promised continuation payoff in a smooth way. If actions have been high in the past, then incentive compatibility requires W to drop to a lower level, which in turn lowers future actions through this term. The second component responds one-for-one to the changes in the target. The first two components add up to the solution in the complete information benchmark, but the third component is unique to the private information case. This new term stipulates how the policy adjusts to state changes, *besides* the “first-best” response in the second component. It is worth noting that while the third term appears because of incentive constraints, it is not for the purpose of keeping incentive compatibility. Indeed, incentives are guaranteed by the diffusion-less evolution of \mathbf{W} , whereas this third component exists for optimality reasons under the presence of the agency problem. In other words, a policy can stipulate alternative reactions to state changes: high current action but lower future actions, or vice versa. Both deliver the same expected payoff to the agent, but different costs to the principal.

Rearranging terms in the policy function, one arrives at the Euler’s equation:

$$\underbrace{x^*(\theta, W) - f(\theta)}_{\text{current distortion}} = \underbrace{W - \gamma \star f(\theta)}_{\text{future distortion}}, \quad (18)$$

which has clear economic interpretations. The left-hand side is the gap between the current action x and the current target $f(\theta)$, and the right-hand side is the gap between the expected future action W and the expected future target $\gamma \star f(\theta)$ given the current state θ (recall the

interpretation of the convolution). In other words, the *current* distortion must balance the *future* distortions for optimality. Given total expected distortions, it is optimal to set the current distortion as a proper share. This smoothing motive is built in the convexity of cost in actions.

Two expressions of the optimal contract are displayed in Proposition 3. The optimal contract (16) directly expresses the dependence of the current action on the entire history. The integral term, $\frac{\sigma^2}{2} \int_0^t (\gamma \star f)''(\theta_s) ds$, is where the report history matters. The history itself is not payoff-relevant, yet it enters to account for the evolution of continuation payoff. The alternative expression (17) shows the contract as a diffusion process with drift $\frac{1}{2} f''(\theta_t)$ and volatility $[f'(\theta_t) - (\gamma \star f)'(\theta_t)] \sigma$. This diffusion representation holds only when f is twice-differentiable everywhere.

5.2.1 Scope of Communication

The optimal contract has its implications about the efficiency of communication in the perspective of the principal. The cost function must be sandwiched by two obvious bounds: the cost \underline{C} from complete information as a lower bound and the cost \overline{C} from babbling as an upper bound. When it coincides with the upper bound, communication is rendered entirely futile by the agent's incentive. On the contrary, when it reaches the lower bound, the agency cost is non-existent. Can valuable information be elicited from the agent without friction? Or is incentive provision so costly that the value of information is totally dissipated? In fact, the shape of the target function is the key to answering these questions.

The comparison with the complete information case is straightforward. If the target is a constant, then the principal does not care about the state and hence there is no cost from the agency problem. Conversely, if the principal obtains the lower bound of cost, then action always matches the target, which will be incentive compatible only if the target is a constant. Proposition 4 gives a formal statement.

Proposition 4 *The cost function coincides with \underline{C} if and only if $f(\cdot)$ is a constant almost everywhere.*

Proof. See Appendix. ■

The comparison to babbling is more complicated. As mentioned at the end of Section 4.2, the disagreement of interest undermines meaningful communication, but the shape of the target function may grant the principal some leeway to reallocate current and future actions so as to salvage some information value. Proposition 3 shows the generic responsiveness of action to information. In the knife-edge case, however, where the volatility of action is always zero, then changes in the state are not reflected in the actions, which is tantamount

to babbling. The next result gives a necessary and sufficient condition for communication to fail.

Theorem 2 (Impossibility)

The cost function coincides with \bar{C} if and only if for some $c_0, c_1, c_2 \in \mathbb{R}$, $f(\theta) = c_0 + c_1\theta + c_2\theta^2$ almost everywhere.

Proof. See Appendix. ■

According to the theorem, if the target function is linear—say, the identity function—then babbling is the inevitable outcome, even with commitment power. More surprisingly, the optimal contract still leads to babbling even if the target is quadratic. That is, the curvature of f does not help facilitate meaningful communication. Theorem 2 conveys the message that it is the curvature of f' , rather than that of f , that helps the principal.

Why is it the curvature of f' , i.e., the third derivative of f , that matters? Here I provide an heuristic explanation. The slope $f'(\theta)$ can be interpreted as the *information sensitivity* of the principal at state θ . The steeper the slope, the more sensitive the target, with respect to a state change. A flatter slope corresponds to a lower information sensitivity. When the slope f' has curvature, the expected information sensitivity in the future $(\gamma \star f)'$ differs from the current f' . This difference leaves the principal with potential gains from reallocating actions between present and future. This is how a dynamic contract facilitates communication above the babbling level. The next section expands this idea to study the direction of trade-off.

5.2.2 Pro- or Counter-Cyclical Policy?

With the optimal contract at hand, it is convenient to analyze its property. In particular, one can now proceed to answer the question raised in the Introduction: whether or not the action necessarily moves in the direction that echos the change of target? Namely, should the principal take high actions in response to high targets? The following analysis gives a negative answer, and provides necessary and sufficient conditions for the action to behave counter-intuitively. To proceed, I first define two terms.

Definition 1 (Cyclicity)

The action x is called pro-cyclical (counter-cyclical, resp.) at state θ if:

$$\text{sgn} \left(\frac{dx}{d\theta} \right) \text{sgn} \left(\frac{df(\theta)}{d\theta} \right) > 0 \text{ (} < 0, \text{ resp.)}$$

In other words, the action is pro-cyclical if it moves in the same direction as the target; otherwise, counter-cyclical.

At any state θ_t such that $f''(\theta_t)$ exists, apply Ito's lemma to (16) to reach its differential form:

$$dx_t^* = \frac{\sigma^2}{2} f''(\theta_t) dt + \sigma [f'(\theta_t) - (\gamma \star f)'(\theta_t)] dZ_t, \quad (19)$$

which has the same drift as, but different volatilities from, the complete information solution:

$$d\underline{x}_t = \frac{\sigma^2}{2} f''(\theta_t) dt + \sigma f'(\theta_t) dZ_t.$$

While the same drift reflects the principal's attempt to match actions with targets *on average*, the difference in volatilities manifests the impact of incentive constraints. Because of the $(\gamma \star f)'(\theta_t)$ term in (19), the sign of the volatility can be reversed, which qualitatively alters the way the action responds to information. In particular, following a positive shock $d\theta_t = \sigma dZ_t > 0$ in the state, the complete information action \underline{x}_t responds by $f'(\theta_t) d\theta_t$ (which is identical to the change in the target), while the action x_t^* in the optimal contract changes by $[f'(\theta_t) - (\gamma \star f)'(\theta_t)] d\theta_t$. The additional convolution term can sometimes reverse the sign of volatility, and when that happens, the action moves counter-cyclically with the target. For example, let $f(\theta) = -e^{-a\theta}$ for some $a \in (0, \sqrt{\frac{r}{2\sigma^2}})$, then $(\gamma \star f)'(\theta) = \frac{2r}{2r - a^2\sigma^2} f'(\theta) > f'(\theta)$ for *all* $\theta \in \mathbb{R}$, hence the action is always counter-cyclical. Figure 3 displays sample paths of state and action for this exponential target function. The dotted curve represents the hypothetical contract in which the action does not have any volatility. The action path and the target path, plotted in solid and dashed curves respectively, always lie on opposite sides of the dotted curve. Below I give a necessary and sufficient condition for the optimal contract to move pro- or counter-cyclically with the target.

Theorem 3 (Conditions for Cyclicity)

The following statements are equivalent:

- (i) *The optimal contract stipulates a pro-cyclical (counter-cyclical, resp.) action at state θ ;*
- (ii) *$\text{sgn}(f'(\theta) - (\gamma \star f)'(\theta)) \text{sgn} f'(\theta) > 0$ (< 0 resp.); and*
- (iii) *$\text{sgn}(\gamma \star f)'''(\theta) \text{sgn} f'(\theta) < 0$ (> 0 resp.).*

Proof. See Appendix. ■

According to the theorem, counter-cyclical actions can occur, and moreover, the condition is not too restrictive. Section 5.4 enumerates some applications where counter-cyclical action is optimal at least for some range of states. At first sight it may seem surprising for a counter-cyclical action to be part of an optimal contract, especially when it is *always* counter-cyclical in Figure 3. After all, it is tempting to increase the action following a rise in the target, the




Figure 3: Sample paths of actions. The parameters are: $r = 1, \sigma = 1, f(\theta) = -e^{-\theta/10}$. The dashed curve is the complete-information action path. The solid curve is the action path in the optimal contract. The dotted curve is the hypothetical evolution of x_t when the volatility term is eliminated; it always separates the target and the action on different sides.

way that the “quota” of actions is usually used in the literature. How can a counter-cyclical action be optimal?

Here are some intuitions. For simplicity, assume $f'(\cdot) > 0$. If the current state increases so that the target is higher, the principal is obviously tempted to increase the current action. Meanwhile, by persistence of the state process, a higher current state predicts higher future states as well, also creating the motive to increase future actions to better match the on-average higher targets. However, she cannot achieve both due to the incentive constraint:

$$\begin{aligned} dW_t &= r(W_t - x_t)dt \\ \iff rx_t dt + dW_t &= rW_t dt, \end{aligned}$$

where dt is used to denote the “current instant” for ease of exposition. The second line is obtained by simply rewriting the first, and it has a similar “incentive budget” interpretation as in the two-period model. Since the agent’s total payoff W_t (inclusive of the flow payoff at dt) must not respond to shocks in the state, the right-hand side is the fixed budget. The principal, facing this budget, can respond to an increase in the state by taking a higher action x_t during dt , in exchange for a lower dW_t (i.e., a faster decline of W_t in the next moment) which translates into a lower sequence of future actions. Or, she can do the opposite, but the incentive budget prohibits her from raising current and future actions simultaneously.

The slope $f'(\theta)$ is the *information sensitivity* of the target at state θ , as mentioned before. If $(\gamma \star f)'''(\theta_t) > 0$, then through the similar logic of Jensen’s inequality, future information sensitivity is greater in expectation than the current one. In other words, following an

increase in the state, the expected future target increases by more, on average, than the current target, or equivalently, the expected future target increases *relative to* the current target. In order to rebalance distortions between the present and future within the incentive budget, the principal optimally lowers the current action in order to better match the future targets, even though it entails a counter-cyclical action against her temptation. The trade-off is flipped when $(\gamma \star f)'''(\theta_t) < 0$, and the action is pro-cyclical. 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 $(\gamma \star f)'''(\theta_t) = 0$, there is no direction for profitable trade-off, and the principal is stuck with babbling.

In the literature of allocation problems (Jackson and Sonnenschein (2007), Renault, Solan, and Vieille (2013), Guo and Hörner (2017)), some version of “quota mechanism” arises: the total allowance or allocation is limited so that the agent must report prudently. In my model, the promised continuation payoff works as a quota as well. The difference lies in *how* the quota is used. Using my terminology, the quota is spent pro-cyclically in the literature: as long as the quota is not depleted, spend it when the state is worth spending, and save it otherwise. In this paper, however, the optimal action can be counter-cyclical: save the quota when the target is high, but spend it when it is low.

To quantify the magnitude of cyclicity, I define the ratio:

$$\xi(\theta) \equiv \frac{f'(\theta) - \gamma \star f'(\theta)}{f'(\theta)}$$

for all θ and call it the *responsiveness factor*. Clearly, pro- (counter-) cyclical action is equivalent to $\xi(\theta)$ being positive (negative) at state θ . Moreover, $\xi \equiv 0$ corresponds to babbling because the volatility of \mathbf{x} is zero and the drift is a constant. For a strictly increasing target f , the range of the responsiveness factor is summarized in Corollary 1.

Corollary 1 *If f is strictly increasing, then $\xi \in (-\infty, 1)$. In particular, the optimal contract never over-shoots.*

The no-over-shooting result indicates that the agency problem always distorts the contract in the direction of less response (or even negative response) to state changes. This is intuitive. Suppose at some state the action responds more than one-for-one, then dampening it to exactly one-for-one not only provides a better matching at the moment, but also alleviates future distortions.

5.2.3 Implementation

Implementation of the optimal contract can be done by dynamically delegating actions to the agent. The principal assigns the agent an “account” of actions. When action x_t is taken at time t , the amount x_t is deducted from the account. At the same time, the account grows at an interest rate r . In this way, the account plays the role of \mathbf{W} , and the agent is indifferent among all action paths. Among them, the principal recommends that the agent act exactly according to (16), and the agent is willing to obey.

To return to the original story in the Introduction, the delegation is in the form of an initially allotted pool of resources that serves as the quota for actions. The division manager does not need to report to the headquarters; instead, he uses resources over time at his own discretion, and the remaining resources grow at the interest rate. The policy function takes the form of a brochure recommending to the manager how much resource to use based on the current profitability and the amount of remaining resources. The brochure may ask the manager to use resources pro- or counter-cyclically depending on the shape of the target function. In Section 5.4, I apply the results to some examples with economically meaningful target functions and derive more implications.

5.3 Evolution of the Contract and Comparative Statics

On path of the optimal contract, the cost and the payoff evolve over time. Also, the property of the optimal contract varies with parameters. In Proposition 5 below, I explore the asymptotics to answer questions such as whether or not the principal faces higher and higher costs over time, and whether or not the agent receives less and less continuation payoff as the relationship evolves. Also, I briefly examine what the contract looks like when the players become very patient/impatient, or when the state process becomes very volatile/stable. The effect on the contract of changing the target function is also discussed.

Proposition 5 (Cost and Payoff Dynamics)

- (i) *The continuation cost is a sub-martingale, i.e., $\frac{\mathbb{E}_t[dC_t^*]}{dt} \geq 0$;*
- (ii) *the continuation payoff monotonically increases (decreases) over time if $\gamma \star f$ is convex (concave). It diverges to ∞ ($-\infty$) if $\gamma \star f$ is strongly convex (concave).*

Proof. See Appendix. ■

The first half of Proposition 5 claims that the principal faces higher and higher costs in expectation as the contract is carried out over time. One may attribute the increasing cost to the impatient principal’s motive to back-load distortions into the future, but that is not the case here. The reason is that with incentive constraints, the promised continuation payoff

diverges away from the principal's favorite level $\gamma \star f$ almost surely, therefore the distortion from promise keeping accumulates over time. It is shown in the proof that the drift of C^* equals $r(C_t^* - \underline{C}_t)$, implying that the drift is strictly positive unless the cost of complete information is obtained, namely f is constant.

The second half predicts the drift of the continuation payoff in two cases, based solely on the curvature of $\gamma \star f$, or equivalently, on the difference between $\gamma \star f$ and f . If f has both convex and concave segments, then W_t may diverge to either ∞ or $-\infty$ without monotonicity, or even oscillate indefinitely. This part of the proposition implies that the agent does not end up immiserated; instead, his destiny depends on the nature of the target function.

There are three parameters to study for comparative statics: the discount rate r ; the volatility σ ; and the target $f(\cdot)$. As $r \rightarrow \infty$, both players become myopic. As a result, the convolution becomes an identity operator and the policy approaches $x_t(W, \theta) = W$, just serving to pay out the promised continuation payoff right away. Iterating this policy function and optimally choosing $W_0 = f(\theta_0)$, the action path tends to the constant $x_t = f(\theta_0)$, which is the babbling outcome for any target function f . The difference between \overline{C} and \underline{C} vanishes, and C^* is sandwiched inbetween. The convergence to the cost of complete information \underline{C} does not mean that communication is effective; it only reflects the fact that future cost weighs little. On the other hand, as $r \rightarrow 0$, players are perfectly patient. The convolution aggregates the target from faraway states with an equal weight as the current state, yielding the expectation of f over an improper distribution. At the same time, the coefficient $\frac{\sigma^2}{r}$ in (14) explodes to infinity because the unavoidable and potentially growing future inefficiency weighs in. As a result, the exploding coefficient dominates and the complete information cost \underline{C} is unobtainable as $r \rightarrow 0$. That being said, the policy approaches the complete information policy if $\lim_{\theta \rightarrow -\infty} f'(\theta) = \lim_{\theta \rightarrow \infty} f'(\theta) = 0$.

The effect of changing σ is, in general, the opposite of changing r , but not exactly. When $\sigma \rightarrow 0$, the convolution again becomes the identity operator and the policy approaches the babbling solution. Nevertheless, in the current setting, the initial θ_0 is commonly known, so this knowledge perfectly predicts future states as well. Due to this assumption, the babbling outcome coincides with the complete information outcome. When $\sigma \rightarrow \infty$, the cost does not converge to the complete information case for the same reason as $r \rightarrow 0$.

The comparative statics on target f are more complicated. As a first step, consider the effect of an affine transformation of the target: $\hat{f}(\theta) \equiv b_0 + b_1 f(\theta)$, which includes special

cases of vertical shift ($b_1 = 0$) and pure scaling ($b_0 = 0$). As is evident from (16) and (14):

$$\hat{x}_t^* = b_0 + b_1(f(\theta_t) + \gamma \star f(\theta_0) - \gamma \star f(\theta_t)) + b_1 \frac{\sigma^2}{2} \int_0^t (\gamma \star f)''(\theta_s) ds = b_0 + b_1 x_t^*,$$

$$\min_{W_0} \hat{C}^*(\theta_0, W_0) = \frac{b_1^2 \sigma^2}{r} \gamma \star (\gamma \star f')^2(\theta_0) = b_1^2 \min_{W_0} C^*(\theta_0, W_0).$$

As can be seen, the constant b_0 serves to shift the entire path of actions without affecting the cost function. The coefficient b_1 scales the action path linearly and the cost quadratically.

Next, let us alter the slope of f by superimposing a linear function around θ_0 , and the target becomes $f(\theta) + b(\theta - \theta_0)$. As a result, x_t does not change at all according to (17), which means that the optimal policy is invariant to rotations in the target. The cost, on the other hand, is affected by this rotation. If f is increasing, then by adding a positive slope to it, the new target is more information-sensitive everywhere. By (14), the cost strictly increases because the optimal policy remains the same, that is, it does not respond to the increased slope.

The third step is to add a non-negative and convex function to the target: $\hat{f}(\theta) \equiv f(\theta) + h(\theta)$ where $h \geq 0$ and $h'' \geq 0$. It turns out that:

$$\hat{x}_t^* - x_t^* = h(\theta_0) + \int_0^t \frac{\sigma^2}{2} h''(\theta_s) ds + \int_0^t (h'(\theta_s) - (\gamma \star h)'(\theta_s)) d\theta_s.$$

If h is quadratic, then $h' - \gamma \star h' \equiv 0$, and $\hat{x}_t^* - x_t^* \geq 0$ for every realization of $(Z_t)_{t \geq 0}$ and every $t \geq 0$. However, if h is not quadratic, then the pointwise comparison is not true. What is still comparable is the distribution of the actions at every time $t \geq 0$. In fact, $\hat{x}_t^* - x_t^* - \int_0^t \frac{\sigma^2}{2} h''(\theta_s) ds$ is a martingale, so that:

$$\mathbb{E}_0 [\hat{x}_t^*] = h(\theta_0) + \mathbb{E}_0 [x_0^*] + \mathbb{E}_0 \left[\int_0^t \frac{\sigma^2}{2} h''(\theta_s) ds \right] \geq \mathbb{E}_0 [x_t^*].$$

Finally, if h is non-negative but not convex, then the comparison fails even in distribution.

5.4 Applications

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

5.4.1 Exponential Target

Exponential target function can be a good approximation if the target increases with the state, but displays increasing or decreasing sensitivity to state changes. In concrete terms, suppose the private state $\tilde{\theta}$ is the profitability or the market's willingness to pay in a product of a division, which follows a geometric Brownian motion:

$$d\tilde{\theta}_t = \frac{\sigma^2}{2}\tilde{\theta}_t dt + \sigma\hat{\theta}_t dZ_t.$$

If we redefine $\theta \equiv \log \hat{\theta}$ as a new state, then θ follows a Brownian motion as in the main model: $d\theta_t = \sigma dZ_t$. If the target action is to match the profitability, then the target can be expressed as an exponential function of the new state:

$$f(\theta_t) = \tilde{\theta}_t = e^{\theta_t}.$$

More generally, an exponential target function can be written as $f(\theta) = be^{\alpha\theta}$ where $|\alpha| < \sqrt{\frac{r}{2\sigma^2}}$ and $b \neq 0$. The bounds on α are required to satisfy Assumption 1. The target is increasing if $b\alpha > 0$ and decreasing if $b\alpha < 0$.

With this exponential target function, the expected future marginal target becomes:

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

Since $f'(\theta) - (\gamma \star f)'(\theta) = -\frac{\alpha^2\sigma^2}{2r - \alpha^2\sigma^2} f'(\theta)$, Theorem 3 predicts that the action always moves counter-cyclically with the state. In fact, the responsiveness factor is a negative constant:

$$\xi_t = \frac{f'(\theta) - (\gamma \star f)'(\theta)}{f'(\theta)} = -\frac{\alpha^2\sigma^2}{2r - \alpha^2\sigma^2} < 0.$$

Note that counter-cyclicity occurs regardless of the signs of b and α .

Returning to the context of the headquarter-manager relationship, the optimal contract is implemented by dynamic delegation with an account for resources, as previously mentioned. Specifically, the headquarter suggests that the manager always takes the counter-cyclical usage of resources. At any time, the manager first discovers the change in the target (the ideal usage of resources), then withholds the temptation to match that change but instead uses resources with the amount against the direction of the change. As a result, resource usage is delayed when profitability is high because its drift is expected to be even higher; resource is brought forth from the future when profitability is low because future drift is even lower on average.

To conclude this example, Proposition 6 finds conditions for the optimal contract to display a negative constant responsiveness factor.

Proposition 6 *The responsiveness factor is a negative constant if and only if $f(\theta) = b_0 + b_1 e^{-\alpha\theta} + b_2 e^{\alpha\theta}$ for some b_0, b_1, b_2 and some $\alpha \in (0, \sqrt{\frac{r}{2\sigma^2}})$.*

Proof. See Appendix. ■

5.4.2 Capped Target

In some situations, the target action has a piecewise nature: initially the state enters the target function linearly without decreasing returns to scale, until the state reaches some threshold level above which the target increases at a lower slope with the state, or even stops growing. Market saturation is an example of this regime change. As another example, the marginal return for a technical parameter drops when a certain threshold is met. Of course, the regime shift can also occur on the lower end of the target function, or on both sides.

Here, I study a case where the target function is capped from both upper and lower ends. Suppose:

$$f(\theta) = \begin{cases} \theta_a & \text{if } \theta < \theta_a \\ \theta & \text{if } \theta_a \leq \theta \leq \theta_b \\ \theta_b & \text{if } \theta > \theta_b \end{cases},$$

where $\theta_a < \theta < b$ are the two bounds. Algebra shows that $f'(\theta) < (\gamma \star f)'(\theta)$ whenever $\theta \notin (\theta_a, \theta_b)$, which is depicted in Figure 4 with parameters $\theta_a = -1$ and $\theta_b = 1$. The left panel shows the target (solid curve) and the expected future target (dashed curve), with the latter steeper than the former only if $\theta \notin (\theta_a, \theta_b)$. The right panel plots the difference in the slopes directly. Theorem 3 predicts the action to move pro-cyclically within (θ_a, θ_b) , but cyclicity is not defined outside that region since $f' = 0$. Nevertheless, the action and the *state* move in opposite directions when the state is in the flat regions. As $\theta_a \rightarrow -\infty$, the target monotonically converges to $\min\{\theta, \theta_b\}$, bounded only from above. As $\theta_b \rightarrow \infty$, only the lower bound θ_a remains. Hence, one-sided bounds are included as special cases by taking the limits.

Intuitively, when the target moves one-for-one with the state, the action should respond pro-cyclically because the information sensitivity f' is already at its highest possible value, and the expected future slope cannot be even higher. However, when the target is capped at the top or bottom, the current information sensitivity is at its lowest value, zero, so that the current action moves against the state to reduce future distortions.

In the headquarters-manager context, the implication for such target functions is clear: allocate resources in phase with the reported technical parameter as long as the ideal amount




Figure 4: Target bounded between -1 and 1. Left panel: f (solid) and $\gamma \star f$ (dashed). Right panel: difference between f' and $(\gamma \star f)'$.

of resources still depends on that parameter, but withhold and delay resource allocation when the ideal amount of resources levels off.

It is important to notice that when $\theta_a = -\infty$, the target function is increasing and concave, which is also true for the exponential target $f(\theta) = -e^{-x}$. For the purpose of modeling a real economic situation, both are good representations of an increasing and concave target, but as we see, the implications for the optimal contract are very different. This corroborates the findings of the paper that it is the third derivative, rather than the second derivative, of the target that matters for the pattern of the contract.

5.4.3 Binary Target

In some applications, the target takes binary values. As a simple example, the product of the division creates revenue for the company if and only if a technical parameter falls into some range, therefore $f(\theta) = \mathbb{1}_{\{\theta \in [\theta_a, \theta_b]\}}$ for some $-\infty \leq \theta_a < \theta_b \leq \infty$. In other words, the headquarters has an “active zone” $[\theta_a, \theta_b]$ in which the target is one, otherwise zero.

While f is discontinuous at θ_a and θ_b , $\gamma \star f$ is continuously differentiable. Figure 5 pictures the target and expected future target, as well as the difference in their slopes. Algebra shows that $f' < (\gamma \star f)'$ when $\theta < \frac{\theta_a + \theta_b}{2}$ and $\theta \neq \theta_a$. Cyclicality cannot be defined anywhere (again, because $f' = 0$ or does not exist), but it can be shown that for states below the midpoint of the active zone, the action moves in the opposite direction of the *state* except at θ_a . What happens when the state crosses θ_a from below? At that moment f jumps up by $f(\theta_{a+}) - f(\theta_{a-}) = 1$ but $\gamma \star f$ is continuous and twice differentiable, hence x jumps up by the same amount as in f , according to (16). Loosely speaking, one can say that the action is pro-cyclical at θ_a with responsiveness factor 1. For states above the midpoint, the contract works like a mirror image. Special cases obtain as $\theta_a \rightarrow -\infty$ or $\theta_b \rightarrow \infty$, and the target becomes a step function with one threshold.




Figure 5: Target is active between -1 and 1. Left panel: f (solid) and $\gamma \star f$ (dashed). Right panel: difference between f' and $\gamma \star f'$.

The implication for the headquarters' resource allocation rule is simple: when the parameter is low (below the midpoint), allocate resources against the movement of the parameter unless the parameter enters or exits the active zone, in which case the flow of resource usage jumps up or down along with the ideal amount. When the parameter is high (above the midpoint), resource allocation and the parameter move in the same direction unless the parameter is at the boundary.

6 Extensions

This section explores the consequences of modifying the main model in various directions.

6.1 The Role of Persistence

Persistence of the state process has demonstrated its importance in driving actions away from tracking the target perfectly. I assume the state process to be a Brownian motion so that any shock persists through time without decay. In the literature, however, a commonly used state process is the finite-state Markov chain, where mean-reversion is automatically built in. To examine the effect of different degrees of persistence on the main results, I modify the state process to be a mean-reversion Brownian motion:

$$d\theta_t = -\mu(\theta_t - \theta_0)dt + \sigma dZ_t.$$

Again, incentive compatibility implies that the continuation payoff W_t does not vary with the current state θ_t , namely, $\beta_t = 0$. With this condition in place of the original set of incentive constraints, the cost of the principal can be expressed recursively as a function of

θ and W :

$$rC(\theta, W) = \min_x r(x - f(\theta))^2 + r(W - x)C_W(\theta, W) - \mu(\theta - \theta_0)C_\theta(\theta, W) + \frac{\sigma^2}{2}C_{\theta\theta}(\theta, W),$$

with the same transversality condition (13).

A closed-form solution is difficult to obtain for a general target function, but for any polynomial target f , the cost function exists in the form of:

$$C^*(\theta, W) = (W - \Gamma_\mu \circ f(\theta))^2 + \frac{\sigma^2}{r}\Gamma_\mu \circ (\Gamma_\mu \circ f)^2(\theta),$$

where the functional operator $\Gamma_\mu \circ f$ is the unique polynomial g of the same order satisfying the ODE:

$$\sigma^2 g''(\theta) - 2(\theta - \theta_0)\mu g'(\theta) - 2rg(\theta) = -2rf(\theta). \quad (20)$$

The policy function has a similar symbolic representation as before:

$$x_t = W_t + f(\theta_t) - \Gamma \circ f(\theta_t),$$

which gives rise to the criterion for counter-cyclical actions:

$$\xi \equiv \frac{f'(\theta_t) - (\Gamma \circ f)'(\theta_t)}{f'(\theta_t)} < 0.$$

Notice that the operator Γ is linear, so in order to obtain the solution for all polynomial targets f , it suffices to solve cases where $f(\theta) = \theta^n$ for all $n \in \mathbb{N}$. As a first step, suppose $n = 0$ so that $f(\theta) = 1$. By observation, $g(\theta) = 1$ is the unique solution to (20) that is a zero-order polynomial. Hence, $\Gamma_\mu \circ f(\theta) = 1 = f(\theta)$.

Next, let $n = 1$ so that $f(\theta) = \theta$ is linear. By the method of undetermined coefficients, it is easy to find the solution:

$$\Gamma \circ f(\theta) = \frac{\mu}{r + \mu}\theta_0 + \frac{r}{r + \mu}\theta.$$

See the left panel of Figure 6 for the comparison between f and $\Gamma \circ f$. The solution $\Gamma \circ f$ is also linear, but the coefficient on θ is *dampened towards zero* by a factor of $\frac{r}{r + \mu}$. Recall that in the main model where $\mu = 0$, this factor is 1. Intuitively, the dampening effect comes from the fact that 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. This effect

has important implications in terms of optimal contract: communication is now meaningful and the action is now pro-cyclical. To see this, note that:

$$\xi(\theta_t) = \frac{dx_t}{df(\theta_t)} = \frac{f'(\theta_t) - (\Gamma \circ f)'(\theta_t)}{f'(\theta_t)} = \frac{\mu}{r + \mu} > 0.$$

As $\mu \rightarrow \infty$, the state process becomes i.i.d., and $\Gamma \circ f$ becomes completely flat at $c_1\theta_0$. The action tracks the target exactly one-for-one, and the cost of the principal achieves the complete information level. The effect of having mean-reversion in the state process explains why communication is meaningful and the quota is used pro-cyclically in the literature with finite state Markov chain.

Further, let $n = 2$ so that $f(\theta) = \theta^2$ is quadratic. A similar method is used to find the unique solution:

$$\Gamma \circ f(\theta) = \frac{2\mu^2\theta_0^2 + (r + \mu)\sigma^2}{(r + \mu)(r + 2\mu)} + \frac{2r\mu\theta_0}{(r + \mu)(r + 2\mu)}\theta + \frac{r}{r + 2\mu}\theta^2.$$

See the right panel of Figure 6 for f and $\Gamma \circ f$. This time, the coefficient for θ^2 is dampened even more by $\frac{r}{r+2\mu}$. With $\mu = 0$, we return to the main model: $\Gamma \circ f(\theta) = f(\theta) + \frac{\sigma^2}{r}$. When $\mu \rightarrow \infty$, the complete information outcome obtains. Just as the case of $n = 1$, the prediction for cyclicity is different from Section 5. The action moves counter-cyclically when θ is between 0 and $\frac{r\theta_0}{2(r+\mu)}$, and the communication is superior to babbling for the principal. This difference, however, is richer than the case $n = 1$. For an interval $\left[0, \frac{r\theta_0}{2(r+\mu)}\right]$ close to the origin, the action becomes counter-cyclical, even though it is not cyclical at all without mean-reversion. Therefore, the dampening effect from mean-reversion does not push actions towards pro-cyclical for *every* state, due to the convolution-like operator which also summarizes the global information of f .

The process can continue for $n > 2$, and the explicit solution can be obtained recursively. As in the quadratic case, it is not guaranteed that stronger mean-reversion leads to more pro-cyclical actions for every state, but in general, pro-cyclicity dominates when the state is far away from θ_0 . Hence, with mean reversion, the counter-cyclicity in the main model is more likely preserved when the state is closer to its balance point; it is more undermined at states further away from the balance point. Proposition 7 formally states the effect of mean reversion for polynomial targets.

Proposition 7 *Cost function and optimal policy exists for all polynomial targets when $\mu > 0$. Moreover, fixing a $\mu > 0$, there exists $\bar{\theta}(\mu) > 0$ s.t. the optimal policy is pro-cyclical whenever $|\theta| > \bar{\theta}(\mu)$.*




Figure 6: The effect of mean-reversion with parameters $\mu = .5, r = 1, \sigma = 1, \theta_0 = 1$. Solid curves are target functions. Dashed curves are expected future target functions. Left panel: $f(\theta) = \theta$. Right panel: $f(\theta) = \theta^2$.

Proof. See Appendix. ■

6.2 Transfer and Limited Liability

It is useful to see how much the main results extend to situations where monetary transfer is either legal or difficult to detect or prohibit. If the transfer can move in either direction, then the socially efficient action is always taken, and the analysis of policy is trivial. More realistically, money only moves from the principal to the agent, i.e., the agent has a limited liability constraint at zero.

For the purpose of emphasizing the role of money, I focus on the simple target function: $f(\theta) = \theta$. The linearity does not result in babbling because there is now transfer. Let $\mathbf{y} = (y_t)_{t \geq 0}$ denote the process of money transfer from the principal to the agent, as part of the contract. Limited liability requires $y_t \geq 0$ for all t . Both players are risk-neutral so that the principal solves the following problem:

$$\begin{aligned} \min_{(x_t)_{t \geq 0}, (y_t \geq 0)_{t \geq 0}} \quad & \mathbb{E}^{\mathbf{l}^\dagger} \left[\int_0^\infty r e^{-rt} ((x_t - f(\theta_t))^2 + y_t) dt \right] \\ \text{s.t.} \quad & \mathbb{E}^{\mathbf{l}^\dagger} \left[\int_0^\infty r e^{-rt} (x_t + y_t) dt \right] \geq \mathbb{E}^{\mathbf{l}} \left[\int_0^\infty r e^{-rt} (x_t + y_t) dt \right], \quad \forall \mathbf{l} \in \mathcal{L}. \end{aligned}$$

The necessary condition for IC is similar: continuation payoff has no volatility. Hence, $dW_t = r(W_t - x_t - y_t)dt$. The homogeneity of f brings the number of state variables down to one, which is $\rho = \theta - W$, the distance between the current state and the promised future actions. In other words, the state variable measures how much the continuation falls short of its expected optimal level, and a ρ too high or too low is costly for the principal. We have

the following functional equation:

$$rC(\rho) = \min_{x-W, y \geq 0} r(x - W - \rho)^2 + ry + r(x - W + y)C'(\rho) + \frac{\sigma^2}{2}C''(\rho). \quad (21)$$

It can be shown that in the optimal contract, transfer is never *actually* used in finite time. Suppose it is used at 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. The fact that money is not used in finite time does not mean that the availability of money does not affect the optimal contract. In fact, money always serves as an *option* that can be used to fulfill the continuation payoff W . If ρ is very low, i.e., the continuation payoff is much higher than the current state, then the principal is tempted to use money rather than high actions to fulfill the promise, since money has a linear cost while action has a quadratic cost. In the limit as $\rho \rightarrow -\infty$, the cost function converges to the one associated with efficient allocation where the use of money is unlimited. The story is different when ρ is very high. The principal is willing to “charge” money from the agent, instead of taking low actions, to fulfill the low continuation payoff, but due to limited liability she cannot do so. Hence, as $\rho \rightarrow \infty$, the cost function converges to the no-transfer case in the main model.

With the intuition above, transversality conditions in both directions can be derived as follows. On the one hand, from (14), the no-transfer cost function for $f(\theta) = \theta$ is $C^*(\theta, W) = (W - \theta)^2 + \frac{\sigma^2}{r} = \rho^2 + \frac{\sigma^2}{r}$. On the other hand, the efficient allocation with unlimited transfer is obtained by taking the first order conditions for x and y in (21). As a result, $C'(\rho) = -1$ and $x = \rho + \frac{1}{2}$, hence the cost function with unlimited transfer is $-\rho - \frac{1}{4}$. In summary, the transversality conditions read:

$$\lim_{\rho \rightarrow \infty} \left(C(\rho) - \rho^2 - \frac{\sigma^2}{r} \right) = 0, \quad \lim_{\rho \rightarrow -\infty} \left(C(\rho) + \rho + \frac{1}{4} \right) = 0.$$

The solution of the cost function is depicted in the left panel of Figure 7. Not surprisingly, it lies between the no-transfer cost and the unlimited-transfer cost, and converges to either of these two bounds when ρ becomes extreme. The policy function is also obtained through $x(\theta, W) = \theta - \frac{1}{2}C'(\theta - W)$. Notice that $x(\theta, W)$ no longer linearly depends on W , meaning that the cyclicity (responsiveness factor) depends not only on the state, but also on history. The middle panel of Figure 7 pictures the responsiveness factor ξ as a function of θ when holding $W = 0$. Consistent with the intuition above, ξ approaches 1, the level in efficient allocation, when θ is low; ξ converges to 0, the level in babbling according to the main model, when θ is high.

Similar result extends to any nonlinear target function $f(\cdot)$, but then the problem cannot




Figure 7: The effect of transfer. Left panel: cost function in solid curve, with $f(\theta) = \theta$. Middle panel: the responsiveness factor (at $W = 0$) in solid curve, with $f(\theta) = \theta$. Right panel: the responsiveness factor (at $W = 0$) in solid curve, with $f(\theta) = -e^{-\theta}$.

be simplified to use only one state variable. Again, money is never actually used in finite time, serving only as an option. Given a state, if the continuation payoff is high enough, then the cost function behaves like the one from unlimited transfer: $W - \gamma \star f(\theta) - \frac{1}{4}$. If the continuation payoff is low, then the cost function is approximately $C^*(\theta, W)$, as seen in the main model. The right panel of Figure 7 shows the responsiveness factor as a function of the state while holding $W = 0$, for the target function $f(\theta) = -e^{-\theta}$. Again, it behaves as in the efficient allocation if the state is low, but the prediction of counter-cyclical action ($\xi < 0$) is preserved when the state is high.

Therefore, with money and limited liability, one can reasonably conjecture that the cyclicity result carries over when the state is high. Pro-cyclical action obtains to track the target one-for-one when the state is low.

6.3 Participation Constraint

In the main model, the agent has no participation constraint. This is for convenience only, as ex ante participation constraint of the agent can be introduced without changing the qualitative results. Formally, the principal now faces an additional constraint at time 0: $W_0 \geq \underline{W}$ for some outside option \underline{W} .

First, notice that if $\underline{W} \leq \gamma \star f(\theta_0)$, then the optimal contract is exactly the same as in the main model. The agent receives $W_0 = \gamma \star f(\theta_0)$ and the participation constraint is not binding.

Next, if $\underline{W} > \gamma \star f(\theta_0)$, then the previous optimal contract violates the condition. Since the cost function is increasing when $W_0 > \gamma \star f(\theta_0)$, the optimal contract can be found at $W_0 = \underline{W}$, which is the lowest promise that still keeps the agent. This change in W_0 does not affect the inter-temporal trade-offs; the action path is simply a vertically translated version of the original one. As a caveat for this case, $W_0 = \underline{W}$ is optimal, *conditional* on the fact

that the principal proposes the contract to the agent. If \underline{W} is sufficiently high, then the cost of fulfilling the promise is greater than the benefit from communication. Specifically, when:

$$\underline{W} > \gamma \star f(\theta_0) + \sqrt{\overline{C}(\theta_0) - \frac{\sigma^2}{r} \gamma \star ((\gamma \star f)^2)(\theta_0)},$$

the principal refuses to sign the contract with the agent in the beginning.

7 Conclusion

This paper uses 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 notoriously futile even if the principal can commit. In contrast, I show that a dynamic contract salvages partial value of information for the principal in most cases, because of her ability to reallocate distortions across time while fixing the incentives of the agent. Nonetheless, the ability to trade-off inter-temporally disappears if the principal’s favorite action displays constant or linear sensitivity to the state. Moreover, the direction of trade-off between present and future leads to counter-intuitive contracts where the principal takes actions that move in the opposite direction of what she should have done with complete information, despite her temptation to close the gap between the action and the favorite action.

The third derivative, or “prudence”, of target function plays the key role in determining the pattern of the contract and the value of communication. The concept of prudence has been used in the literature of decision under uncertainty, including the precautionary savings motive in consumption theory. This paper shares a similar feature in that the curvature captures the comparison of marginal terms between now and the future, with *marginal utility* appearing in the consumption theory and *information sensitivity* arising in my model. However, with the fixed quota of actions from incentive constraints, the actions can display counter-cyclical to the state, which is not found in precautionary savings.

Moreover, at an abstract level, the counter-cyclical can be thought of as a new implication from agency problems. The seemingly distrusting contract does not directly reflect the suspicion of the principal (recall that it is a truthful contract), but at some fundamental level it indeed stems from conflicts of interest as it appears only when there is agency problem. The more aligned the preferences, the less likely that the optimal contract is counter-cyclical.

The model is built on the simplifying assumption that the principal suffers a quadratic cost from mismatching the action with the target, and the agent’s payoff is linear in action. Although the general logic of the main results extend beyond this specific setting, this

assumption is restrictive. For example, with the quadratic cost in place, the simple form of the state process is no longer a normalization.

References

- ABREU, D., D. PEARCE, AND E. STACCHETTI (1986): “Optimal cartel equilibria with imperfect monitoring,” *Journal of Economic Theory*, 39(1), 251–269.
- ALONSO, R., AND N. MATOUSCHEK (2008): “Optimal delegation,” *The Review of Economic Studies*, 75(1), 259–293.
- AMADOR, M., AND K. BAGWELL (2013): “The theory of optimal delegation with an application to tariff caps,” *Econometrica*, 81(4), 1541–1599.
- ANTIČ, N., AND K. STEVERSON (2016): “Screening through coordination,” *Working paper*.
- AUMANN, R. J., AND S. HART (2003): “Long cheap talk,” *Econometrica*, 71(6), 1619–1660.
- BATTAGLINI, M. (2002): “Multiple referrals and multidimensional cheap talk,” *Econometrica*, 70(4), 1379–1401.
- CHAKRABORTY, A., AND R. HARBAUGH (2010): “Persuasion by cheap talk,” *The American Economic Review*, 100(5), 2361–2382.
- CRAWFORD, V. P., AND J. SOBEL (1982): “Strategic information transmission,” *Econometrica*, 50(6), 1431–1451.
- DEMARZO, P. M., AND Y. SANNIKOV (2016): “Learning, termination, and payout policy in dynamic incentive contracts,” *The Review of Economic Studies*, 84(1), 182–236.
- FERNANDES, A., AND C. PHELAN (2000): “A recursive formulation for repeated agency with history dependence,” *Journal of Economic Theory*, 91(2), 223–247.
- GOLOSOV, M., V. SKRETA, A. TSYVINSKI, AND A. WILSON (2014): “Dynamic strategic information transmission,” *Journal of Economic Theory*, 151, 304–341.
- GOLTSMAN, M., J. HÖRNER, G. PAVLOV, AND F. SQUINTANI (2009): “Mediation, arbitration and negotiation,” *Journal of Economic Theory*, 144(4), 1397 – 1420.
- GREEN, J. R., AND N. L. STOKEY (2007): “A two-person game of information transmission,” *Journal of Economic Theory*, 135(1), 90 – 104.

- GUO, Y., AND J. HÖRNER (2017): “Dynamic mechanisms without money,” *Working paper*.
- HE, Z., B. WEI, J. YU, AND F. GAO (2017): “Optimal long-term contracting with learning,” *The Review of Financial Studies*, 30(6), 2006–2065.
- HOLMSTRÖM, B. R. (1977): “On incentives and control in organizations,” .
- JACKSON, M. O., AND H. F. SONNENSCHNEIN (2007): “Overcoming incentive constraints by linking decisions,” *Econometrica*, 75(1), 241–257.
- KAPIČKA, M. (2013): “Efficient allocations in dynamic private information economies with persistent shocks: a first-order approach,” *Review of Economic Studies*, 80(3), 1027–1054.
- KRISHNA, V., AND J. MORGAN (2001): “A model of expertise,” *The Quarterly Journal of Economics*, 116(2), 747–775.
- (2004): “The art of conversation: eliciting information from experts through multi-stage communication,” *Journal of Economic Theory*, 117(2), 147 – 179.
- RENAULT, J., E. SOLAN, AND N. VIEILLE (2013): “Dynamic sender–receiver games,” *Journal of Economic Theory*, 148(2), 502–534.
- SANNIKOV, Y. (2008): “A continuous-time version of the principal-agent problem,” *The Review of Economic Studies*, 75(3), 957–984.
- THOMAS, J., AND T. WORRALL (1990): “Income fluctuation and asymmetric information: An example of a repeated principal-agent problem,” *Journal of Economic Theory*, 51(2), 367–390.
- WILLIAMS, N. (2011): “Persistent private information,” *Econometrica*, 79(4), 1233–1275.

Appendix

Solving the Two-Period Contract. The IC's are simplified to one equation: $x_1(\theta_1) + \delta x_2(\theta_1) = W$. Plug this back to the objective to yield the unconstrained problem:

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

For every θ_1 , the FOC for $x_1(\theta_1)$ is

$$\begin{aligned} \frac{2}{\delta}(W - x_1(\theta_1)) + 2(x_1(\theta_1) + \mathbb{E}[f(\theta_2)|\theta_1] - f(\theta_1)) &= 0 \\ \Leftrightarrow x &= \frac{W}{1 + \delta} + \frac{\delta}{1 + \delta}(f(\theta_1) - \mathbb{E}[f(\theta_2)|\theta_1]). \end{aligned} \quad (22)$$

The FOC for W reads (after plugging (22))

$$\begin{aligned} \mathbb{E} \left[-\frac{2(-W + \delta \mathbb{E}[f(\theta_2)|\theta_1] + f(\theta_1))}{1 + \delta} \right] &= 0 \\ \Leftrightarrow W &= \mathbb{E}f(\theta_1) + \delta \mathbb{E}f(\theta_2). \end{aligned}$$

Plugging the above back to (22) and the IC constraint, we have the solution (3) and (4). ■

Proof of Lemma 1. Suppose the arbitrarily given contract \mathbf{x} induces a strategy \mathbf{l} , which may not be truthful. Let $L_t \equiv \int_0^t l_s ds$. Now consider a new contract \mathbf{x}^\dagger such that $x_t^\dagger(\hat{\theta}^t) \equiv x_t(\hat{\theta}^t + L^t)$. I claim that truth-telling \mathbf{l}^\dagger is optimal for the sender under the new contract. If not, then there exists a lying strategy $\mathbf{l}' \neq 0$ along with $L'_t \equiv \int_0^t l'_s ds$ such that $\mathbb{E}^{\mathbf{l}'}$ $\left[\int_0^\infty r e^{-rt} x_t^\dagger dt \right] > \mathbb{E}^{\mathbf{l}^\dagger}$ $\left[\int_0^\infty r e^{-rt} x_t^\dagger dt \right]$. But that contradicts the optimality of \mathbf{l} in the original contract because

$$\begin{aligned} \mathbb{E} \left[\int_0^\infty r e^{-rt} x_t(\theta^t + L^t + L'^t) dt \right] &= \mathbb{E} \left[\int_0^\infty r e^{-rt} x_t^\dagger(\theta^t + L^t) dt \right] \\ &> \mathbb{E} \left[\int_0^\infty r e^{-rt} x_t^\dagger(\theta^t) dt \right] &= \mathbb{E} \left[\int_0^\infty r e^{-rt} x_t(\theta^t + L^t) dt \right]. \end{aligned}$$

The new contract \mathbf{x}^\dagger implements the original mapping from θ^∞ to x^∞ by definition. ■

Proof of Lemma 2. Define the process of total payoff of the sender evaluated at time t with information \mathcal{F}_t :

$$W_0^{\mathcal{F}_t} \equiv \int_0^t r e^{-rs} x_s ds + e^{-rt} W_t,$$

which is a martingale because

$$\mathbb{E}W_0^{\mathcal{F}_t} = \mathbb{E} \left[\int_0^t r e^{-rs} x_s ds \right] + e^{-rt} \mathbb{E} \left[\int_t^\infty r e^{-r(s-t)} x_s ds \right] = W_0^{\mathcal{F}_0} = W_0.$$

By Theorem 1.3.13 in Karatzas and Shreve (1991), the martingale $W_0^{\mathcal{F}_t}$ has a RCLL modification. Therefore by Theorem 3.4.15 in the same book, the martingale has a representation

$$W_0^{\mathcal{F}_t} = W_0^{\mathcal{F}_0} + \int_0^t r e^{-rs} \beta_s \sigma dZ_s, \quad \forall t \geq 0.$$

Subtracting the two expressions for $W_0^{\mathcal{F}_t}$ and then differentiating w.r.t. t , we have

$$dW_t = r(W_t - x_t)dt + r\beta_t \sigma dZ_t = r(W_t - x_t)dt + r\beta_t d\hat{\theta}_t,$$

which has an equivalent integral form

$$W_t = W_0 + \int_0^t r(W_s - x_s)ds + \int_0^t r\beta_s \sigma dZ_s.$$

■

Proof of Proposition 1. With the restriction of strategy space to \mathcal{L} , Novikov's condition is satisfied. By Girsanov Theorem there exists a martingale \mathbf{M} with $M_t \equiv e^{\frac{1}{\sigma} \int_0^t l_s dZ_s - \frac{1}{2\sigma^2} \int_0^t l_s^2 ds}$, serving as the Radon-Nikodym derivative between the measure induced by l and the measure under truth-telling. The martingale evolves according to $dM_t = M_t \frac{l_t}{\sigma} dZ_t$ with $M_0 = 1$. Besides M_t , $L_t = \int_0^t l_s ds$ is also included as a state variable, whose evolution is simply $dL_t = l_t dt$. With these notations, the sender's payoff from an arbitrary strategy \mathbf{l} is $\mathbb{E}^l \left[\int_0^\infty r e^{-rt} x_t dt \right]$, or equivalently $\mathbb{E}^{l^\dagger} \left[\int_0^\infty r e^{-rt} M_t x_t dt \right]$.

Let p^M be the costate variable for the drift of M , and q^M the costate for the volatility of M . Let p^L and q^L be the counterparts for L . The current value Hamiltonian associated with the sender's problem is

$$rMx + q^M \frac{Ml}{\sigma} + p^L l.$$

The first order condition for the control $l = 0$ to be optimal, evaluated at $l = 0, M = 1$, is

$$\frac{q^M}{\sigma} + p^L = 0. \tag{23}$$

The Euler equations for M and L evaluated at $l = 0, M = 1$ are

$$dp^M = r(p^M - x)dt + \frac{q^M}{\sigma}(\sigma dZ_t), \quad (24)$$

$$dp^L = rp^L dt + \frac{q^M}{\sigma}(\sigma dZ_t), \quad (25)$$

with transversality conditions $\lim_{t \rightarrow \infty} p_t^M e^{-rt} = 0$ and $\lim_{t \rightarrow \infty} p_t^L e^{-rt} = 0$. The solution to the above BSDE's are

$$p_t^M = \mathbb{E}_t \left[\int_t^\infty r e^{-r(s-t)} x_s ds \right] = W_t,$$

$$p_t^L = 0,$$

where W_t is the sender's on-path continuation payoff defined in Section 4.2. Hence, by comparing (24) and (8), we have $\frac{q^M}{\sigma} = r\beta$. Plugging this back to (23) and using the fact $p^L = 0$, we have the necessary condition $\beta = 0$. The counterpart of the marginal continuation payoff in the literature is p^L here, which is identically zero. ■

Proof of Proposition 2. For any arbitrary strategy $l \in \mathcal{L}$, define W_t^l and x_t^l as the off-equilibrium continuation payoff and action at time t induced by l . In the beginning, there is no deviation, so $W_0^l = W_0$. Since $\beta_t = 0$, we have $dW_t^l = r(W_t^l - x_t^l)dt$. It follows that

$$d(e^{-rt}W_t^l) = e^{-rt}r(W_t^l - x_t^l)dt - re^{-rt}W_t^l dt = -rx_t^l e^{-rt} dt,$$

hence the agent's expected total payoff from strategy l is:

$$\mathbb{E} \left[\lim_{t \rightarrow \infty} \int_0^t r e^{-rs} x_s^l ds \right] = W_0 - \mathbb{E} \left[\lim_{t \rightarrow \infty} e^{-rt} W_t^l \right].$$

It suffices to show that $\lim_{t \rightarrow \infty} e^{-rt} W_t^l = 0$ so that the actual payoff from any deviation is always W_0 . Let $L_t = \int_0^t l_s ds$ be the cumulative deviation. First, notice that

$$\begin{aligned} & \mathbb{E} [e^{c_1(\theta_t + L_t)} + e^{-c_1(\theta_t + L_t)}] \leq \mathbb{E} [(e^{c_1\theta_t} + e^{-c_1\theta_t})(e^{c_1L_t} + e^{-c_1L_t})] \\ & \leq \sqrt{\mathbb{E} [(e^{-c_1L_t} + e^{c_1L_t})^2]} \sqrt{\mathbb{E} [(e^{-c_1\theta_t} + e^{c_1\theta_t})^2]} \\ & \leq 4\sqrt{2 + 2e^{2c_1^2\sigma^2t} \cosh(2c_1\theta_0)} \sqrt{\mathbb{E} [e^{2c_1|L_t|}]} \\ & \leq 8e^{c_1^2\sigma^2t} \cosh(c_1\theta_0) \sqrt{\mathbb{E} [e^{2c_1|L_t|}]}. \end{aligned}$$

Next, using Hölder's inequality multiple times, we have

$$\begin{aligned}
& \sqrt{\mathbb{E}[e^{2c_1|L_t|}]} \leq \sqrt{\mathbb{E}\left[e^{2c_1\sqrt{\int_0^t l_s^2/\sigma^2 ds}}\sqrt{t}\right]} \\
& \leq \sqrt{\mathbb{E}\left[e^{\sqrt{\frac{1}{8}\int_0^t l_s^2/\sigma^2 ds}}e^{4\sqrt{2}c_1\sqrt{t}}\right]} \\
& \leq \sqrt{\sqrt{\mathbb{E}\left[e^{\sqrt{\frac{1}{2}\int_0^t l_s^2/\sigma^2 ds}}\right]}e^{4\sqrt{2}c_1\sqrt{t}}} \\
& = e^{2\sqrt{2}c_1\sqrt{t}}\sqrt{\sqrt{\mathbb{E}\left[e^{\sqrt{\frac{1}{2}\int_0^t l_s^2/\sigma^2 ds}|A}\right]\mathbb{P}(A)} + \mathbb{E}\left[e^{\sqrt{\frac{1}{2}\int_0^t l_s^2/\sigma^2 ds}|A^c}\right]\mathbb{P}(A^c)}} \\
& \leq e^{2\sqrt{2}c_1\sqrt{t}}e^{rt/2}\left(e^{-2rt}\left(e + \mathbb{E}\left[e^{\frac{1}{2}\int_0^t l_s^2/\sigma^2 ds}\right]\right)\right)^{\frac{1}{4}} \equiv e^{2\sqrt{2}c_1\sqrt{t}}e^{rt/2}H(t),
\end{aligned}$$

where in the equality A represents the event that $\sqrt{\frac{1}{2}\int_0^t l_s^2/\sigma^2 ds} < 1$, and $\lim_{t \rightarrow \infty} H(t) = 0$ by definition of \mathcal{L} . Finally,

$$\begin{aligned}
e^{-rt}\mathbb{E}|W_t^l| &= e^{-rt}\mathbb{E}\left|W_0 + \int_0^t \frac{dW_s}{ds} ds\right| \\
&\leq e^{-rt}\mathbb{E}\left|W_0 + \int_0^t |f(\theta_s) - \gamma \star f(\theta_s)| ds\right| \\
&\leq e^{-rt}|W_0| + c_2 e^{-rt} \int_0^t \mathbb{E}\left[e^{c_1(\theta_s+L_s)} + e^{-c_1(\theta_s+L_s)}\right] ds \\
&= e^{-rt}|W_0| + c_3 e^{-rt} \int_0^t e^{2\sqrt{2}c_1\sqrt{s+rs}/2+c_1^2\sigma^2 s} H(s) ds.
\end{aligned}$$

Since $e^{2\sqrt{2}c_1\sqrt{s+rs}/2+c_1^2\sigma^2 s} H(s)$ is dominated by $e^{\frac{s}{4}(3r+2c_1^2\sigma^2)}$ when t is large enough, $e^{-rt}\mathbb{E}|W_t^l|$ converges to 0 as $t \rightarrow \infty$, and so does $e^{-rt}\mathbb{E}W_t^l$. ■

Solution to Problem (9). Define $z_t \equiv x_t - f(\theta_t) - (W_0 - \mathbb{E}_0[\int_0^\infty re^{-rs}f(\theta_s)ds])$. Then for x to deliver the promised payoff W_0 , we must have $\mathbb{E}_0[\int_0^\infty re^{-rs}z_s ds] = 0$. The receiver's payoff reads

$$\begin{aligned}
& -\mathbb{E}_0\left[\int_0^\infty re^{-rt}\left(W_0 - \mathbb{E}_0\left[\int_0^\infty re^{-rs}f(\theta_s)ds\right] + z_t\right)^2 dt\right] \\
& \leq -\left(W_0 - \mathbb{E}_0\left[\int_0^\infty re^{-rs}f(\theta_s)ds\right] + \mathbb{E}_0\left[\int_0^\infty re^{-rt}z_t dt\right]\right)^2 \\
& = -\left(W_0 - \mathbb{E}_0\left[\int_0^\infty re^{-rs}f(\theta_s)ds\right]\right)^2,
\end{aligned}$$

where the inequality follows from Jensen's Inequality, and equality is obtained when $z_t \equiv 0$. Plugging it back, an optimal solution is $x_t = f(\theta_t) + (W_0 - \mathbb{E}_0 [\int_0^\infty r e^{-rs} f(\theta_s) ds])$. ■

Proof of Theorem 1. The proof requires three steps. First, I show that the candidate cost function indeed achieves the lowest cost in the relaxed problem with policy function (15). Second, I show the uniqueness of the cost function satisfying the HJB and the transversality. Finally, I show that the sufficient condition in Section 4.3 is satisfied.

First, for any contract \hat{x} satisfying the IC condition, define \hat{W} as the resulting process of continuation payoff, and define

$$\hat{C}_t^0 \equiv \int_0^t r e^{-rs} (\hat{x}_s - f(\theta_s))^2 ds + e^{-rt} C^*(\theta_t, \hat{W}_t) \quad (26)$$

as the total cost of the principal discounted at time 0 but evaluated at time t , where the arbitrary contract is used before t but the candidate cost function is obtained afterwards. The goal is to show that \hat{C}_t^0 is a martingale if \hat{x} coincides with the optimal policy (15), and is a sub-martingale if not. The total differential for \hat{C}_t^0 is

$$\begin{aligned} e^{rt} d\hat{C}_t^0 &= r(\hat{x}_t - f(\theta_t))^2 dt - rC^*(\theta_t, \hat{W}_t) dt + r(\hat{W}_t - \hat{x}_t) V_W(\theta_t, \hat{W}_t) dt \\ &\quad + \sigma V_\theta(\theta_t, \hat{W}_t) dZ_t + \frac{\sigma^2}{2} V_{\theta\theta}(\theta_t, \hat{W}_t) dt \\ &= \sigma V_\theta(\theta_t, \hat{W}_t) dZ_t + r(x_t - \hat{x}_t)^2 dt, \end{aligned}$$

where x_t is the optimal policy, and the second equation is obtained by plugging the HJB (12) and the policy function (15). It is clear that

$$\frac{\mathbb{E}_t(e^{rt} d\hat{V}_t^0)}{dt} = r(x_t - \hat{x}_t)^2 \geq 0, \quad (27)$$

with equality only if $\hat{x}_t = x_t$ for all t . Hence, the contract in the proposition is indeed optimal.

The cost function satisfies the transversality condition. Specifically, from Assumption 1, $(\gamma \star f)'$ is exponentially bounded:

$$|(\gamma \star f)'^2(\theta)| = \frac{2r}{\sigma^2} \left| \int_0^\infty \gamma(z) f(\theta - z) dz - \int_{-\infty}^0 \gamma(z) f(\theta - z) dz \right|^2 < \alpha_2 \cosh^2(\alpha_1 \theta),$$

where $\alpha_2 = \frac{16\alpha_0^2 r^3}{\sigma^2(\alpha_1^2 \sigma^2 - 2r)^2} > 0$, so that

$$\sigma^2 \gamma \star ((\gamma \star f)')(\theta) < \frac{\alpha_2 \sigma^2 (r \cosh^2(\alpha_1 \theta) - \alpha_1^2 \sigma^2)}{r - 2\alpha_1^2 \sigma^2}.$$

Now,

$$\frac{\mathbb{E}_t [dC_t^*]}{dt} = \frac{d}{dt} \mathbb{E}_t \left[(W_t - \gamma \star f(\theta_t))^2 + \frac{\sigma^2}{r} \gamma \star ((\gamma \star f)')(\theta_t) \right] = \sigma^2 \gamma \star ((\gamma \star f)')(\theta_t),$$

so that

$$\begin{aligned} |\mathbb{E} [e^{-rt} C_t^*]| &= e^{-rt} \left| C_0^* + \int_0^t \int_{-\infty}^{\infty} \frac{\mathbb{E}_s [dC_s^*]}{ds} d\Phi_s(\theta_s) ds \right| \\ &\leq e^{-rt} |C_0^*| + e^{-rt} \left| \frac{\alpha_2}{4} \left(2\sigma^2 t + \frac{(e^{2\alpha_1^2 \sigma^2 t} - 1) r \cosh(2\alpha_1 \theta_0)}{\alpha_1^2 (r - 2\alpha_1^2 \sigma^2)} \right) \right|. \end{aligned}$$

Taking the limit as $t \rightarrow \infty$ yields the transversality.

Second, to show that any other solution \tilde{C} to (12) and (13) must coincide with C^* , notice that following the logic of (26) and (27),

$$C^* = \hat{C}_0^0 \leq \hat{C}_\infty^0 = \tilde{C},$$

and when we switch \tilde{C} and C^* , it holds that $\tilde{C} \leq C^*$. Consequently, $C^* = \tilde{C}$.

Finally, it remains to check the sufficient condition in Proposition 2. From the policy function (15), we have $x_t - W_t = f(\theta_t) - \gamma \star f(\theta_t)$. From Assumption 1,

$$|f(\theta_t)| \leq \alpha_0 (e^{\alpha_1 \theta} + e^{-\alpha_1 \theta}), \quad |\gamma \star f(\theta_t)| \leq \frac{2\alpha_0}{2r - \alpha_1^2 \sigma^2} (e^{\alpha_1 \theta} + e^{-\alpha_1 \theta}).$$

The bound for $|x_t - W_t|$ is obtained from $|x_t - W_t| \leq |f(\theta_t)| + |\gamma \star f(\theta_t)|$. ■

Proof for Proposition 3. The policy function is obtained by taking FOC of the cost function.

To obtain the optimal contract without explicitly relying on W , derive its evolution as follows

$$\begin{aligned} W_t &= \gamma \star f(\theta_0) + \int_0^t dW_s \\ &= \gamma \star f(\theta_0) + \int_0^t r(\gamma \star f(\theta_s) - f(\theta_s)) ds = \gamma \star f(\theta_0) + \frac{\sigma^2}{2} \int_0^t (\gamma \star f)''(\theta_s) ds. \end{aligned}$$

Hence, plug W_t in (15) to obtain (16). At generic states where f is twice differentiable, apply Ito's lemma to (16) to obtain (17). ■

Proof for Proposition 4. To compare with complete information, notice that since $(\gamma \star f)^2(\theta) \geq 0$ and the convolution preserves the sign, we have $\gamma \star ((\gamma \star f)^2)(\theta) \geq 0$, meaning that $C^* \geq \underline{C}$ for all θ and W . The inequality binds and the complete information cost obtains if and only if $\gamma \star ((\gamma \star f)^2)(\theta) \equiv 0$ almost everywhere, which means $(\gamma \star f)' \equiv 0$ almost everywhere. Notice that because of Assumption 1, $\gamma \star f$ is continuously differentiable, so $(\gamma \star f)' \equiv 0$ holds for all θ . Equivalently, f is a constant almost everywhere. ■

Proof for Theorem 2. For comparison with babbling, notice that babbling outcome is the result of a deterministic action path, hence satisfies IC. By definition, babbling cannot outperform the optimal contract. The following is to establish necessity and sufficiency for them to coincide.

Similar to the proof for Theorem 1, for the babbling action path \hat{x}_t , define \hat{C}_t^0 the same way as in (26). Its drift satisfies $e^{rt} \frac{\mathbb{E}_t[d\hat{C}_t^0]}{dt} = r(x_t - \hat{x}_t)^2$. In order to achieve the babbling cost, we need $x_t = \hat{x}_t$ almost surely w.r.t. the product measure of time and sample space. An immediate implication is that the optimal policy should also be state-independent almost surely. Through (15), this requires $f - \gamma \star f$ to be a constant for almost all θ . Hence, $(\gamma \star f)'' = -\frac{2r}{\sigma^2}(f - \gamma \star f)$ is a constant almost everywhere. From Assumption 1, $\gamma \star f$ is twice differentiable, so that $(\gamma \star f)''$ is always 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.

Conversely, if $f(\theta) = c_0 + c_1\theta + c_2\theta^2$ almost everywhere, then $(\gamma \star f)''$ is a constant, and $f - \gamma \star f$ is almost always a constant. According to (16), $x_t = f(\theta_0) - \frac{\sigma^2}{2}2c_2t$ almost surely. On the other hand, the babbling action \hat{x}_t minimizes

$$\int_{-\infty}^{\infty} r(\hat{x}_t - f(\theta_t))^2 \phi\left(\frac{\theta_t - \theta_0}{\sigma\sqrt{t}}\right) d\theta_t$$

as a quadratic in \hat{x}_t , and the solution coincides with x_t except on zero-measure sets. ■

Proof for Theorem 3. For the first statement, I prove sufficiency first. If f' is convex, then

$$(\gamma \star f)'(\theta) = \int_{-\infty}^{\infty} \gamma(z)f'(\theta + z)dz \geq \int_{-\infty}^{\infty} \gamma(z)(f'(\theta) + zf''(\theta))dz = f'(\theta),$$

where $f''(\theta)$ is a subgradient at θ . Hence, $\frac{dx_t}{dZ_t} = \sigma(f'(\theta_t) - (\gamma \star f')(\theta_t)) \leq 0$ and equality is obtained when f' is strictly convex.

The inequalities above change directions if f' is concave, and then $\frac{dx_t}{dZ_t} = \sigma(f'(\theta_t) - (\gamma \star f')(\theta_t)) \geq 0$.

For necessity, suppose the policy is always counter-cyclical, i.e. $f' < (\gamma \star f)'$. Twice differentiate $(\gamma \star f)'$ to get

$$(\gamma \star f)''' = -\frac{2r}{\sigma^2}(f' - (\gamma \star f)') > 0,$$

hence the convexity of $(\gamma \star f)'$. The above inequality is reversed if the policy is always pro-cyclical.

For the second statement, notice that the argument in the last paragraph does not require global properties of actions to determine the local convexity or concavity in $(\gamma \star f)'$. ■

Proof for Proposition 5. To show (i), note that

$$\begin{aligned} \frac{\mathbb{E}_t [dC_t^*]}{dt} &= \frac{d}{dt} \mathbb{E}_t \left[(W_t - \gamma \star f(\theta_t))^2 + \frac{\sigma^2}{r} \gamma \star ((\gamma \star f)'^2)(\theta_t) \right] \\ &= \sigma^2 \gamma \star ((\gamma \star f)'^2)(\theta_t) \geq 0, \end{aligned}$$

where the second equality follows from Ito's lemma and the policy function.

To show (ii), plug the policy function into (8) and let $\beta_t = 0$ to obtain

$$dW_t = -r(f(\theta_t) - \gamma \star f(\theta_t))dt = \frac{\sigma^2}{2}(\gamma \star f)''(\theta_t)dt,$$

where the second equality comes from the property of γ . ■

Proof for Proposition 6. In order for the responsiveness to be a negative constant $\bar{\xi}$, one has the following equivalence:

$$(\gamma \star f)'(\theta) = (1 - \bar{\xi})f'(\theta).$$

Integrate both sides to obtain $\gamma \star f(\theta) = (1 - \bar{\xi})f(\theta) + b_0\bar{\xi}$ where b_0 is a constant. Differentiating both sides to obtain

$$(1 - \bar{\xi})f''(\theta) = (\gamma \star f)''(\theta) = \frac{2r}{\sigma^2}(\gamma \star f(\theta) - f(\theta)) = \frac{2r}{\sigma^2}\bar{\xi}(b_0 - f(\theta)).$$

Solve it as a second order ODE for f , and one obtains

$$f(\theta) = b_0 + b_1 e^{-\alpha\theta} + b_2 e^{\alpha\theta},$$

where $\alpha = \frac{\sqrt{-2r\xi}}{\sigma\sqrt{1-\xi}}$. ■

Proof for Proposition 7. The cases for $n = 0, 1, 2$ have been proved. For $n \geq 3$, let $f(\theta) = \theta^n$. Conjecture that $\Gamma \circ f(\theta) = \sum_{k=0}^n a_k \theta^k$, and the coefficients are solved below. Plug the conjectured form into the ODE (20), and one obtains

$$2r\theta^n + \sigma^2 \sum_{k=0}^n (k-1)k\theta^{k-2}a_k - 2\mu(\theta - \theta_0) \sum_{k=0}^n k\theta^{k-1}a_k - 2r \sum_{k=0}^n \theta^k a_k = 0.$$

Equating coefficients for θ^n and θ^{n-1} , we have

$$a_n = \frac{r}{r + n\mu}, \quad a_{n-1} = \frac{n\mu\theta_0 a_n}{r + (n-1)\mu}.$$

The other coefficients are obtained by iteration:

$$a_{n-2} = \frac{(n-1)(2\mu\theta_0 a_{n-1} + n\sigma^2 a_n)}{2(r + (n-2)\mu)},$$

until a_0 is solved. Because of the linearity of the ODE (20) and the Γ operator, the solution for $f(\theta) = \sum_{k=0}^n b_k \theta^k$ is obtained by linear combinations: $\Gamma \circ f(\theta) = \sum_{j=0}^n \sum_{k=0}^j b_j a_{jk} \theta^k$.

Following exactly the same approach in the proof for Theorem 1, the policy $x(\theta, W) = W + f(\theta) - \Gamma \circ f(\theta)$ is indeed optimal.

The responsiveness factor is

$$\xi(\theta) = \frac{f'(\theta) - (\Gamma \circ f)'(\theta)}{f'(\theta)}.$$

If f is linear, then $\xi(\theta) = \frac{\mu}{r+\mu}$, and the second statement is true. If f has a highest power of $n \geq 2$, then both the numerator and the denominator diverges to infinity. l'Hôpital's rule applies here:

$$\lim_{|\theta| \rightarrow \infty} \xi(\theta) = \frac{n\mu}{r + n\mu} > 0.$$

Because of the continuity of ξ in θ , there exists $\bar{\theta}(\mu) > 0$ such that $\xi(\theta) > 0$ for all $|\theta| > \bar{\theta}(\mu)$. ■

Relaxing the Strategy Set \mathcal{L} . 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 $T \rightarrow \infty$. At time T , the cost gap between the truncated contract and the optimal one is

$$\begin{aligned} \Delta(\theta_T) &\equiv \gamma \star (W - f(\theta_T))^2 - (W - \gamma \star f(\theta_T))^2 - \frac{\sigma^2}{r} \gamma \star ((\gamma \star f)^2)(\theta_T) \\ &= \gamma \star f^2(\theta_T) - (\gamma \star f(\theta_T))^2 - \frac{\sigma^2}{r} \gamma \star ((\gamma \star f)^2)(\theta_T) \\ &\leq \gamma \star f^2(\theta_T) + (\gamma \star f(\theta_T))^2 + \frac{\sigma^2}{r} \gamma \star ((\gamma \star f)^2)(\theta_T) \\ &\leq b_1 + b_2 \cosh^2(\alpha_1 \theta_T), \end{aligned}$$

where $b_1 = \frac{2\alpha_0^2(12r^2 - 4r\alpha_1^2\sigma^2 + \alpha_1^4\sigma^4)}{(2r - \alpha_1^2\sigma^2)^2} > 0$ and $b_2 = \frac{2r\alpha_0^2(16r^2 - 12r\alpha_1^2\sigma^2 + \alpha_1^4\sigma^4)}{(r - 2\alpha_1^2\sigma^2)(2r - \alpha_1^2\sigma^2)^2} > 0$, and the second inequality follows from Assumption 1. Hence,

$$\begin{aligned} C^T(\theta_0, W_0) - C^*(\theta_0, W_0) &= e^{-rT} \mathbb{E}[\Delta(\theta_T)] \\ &< e^{-rT} \left(b_1 + \frac{1}{2}b_2 + \frac{1}{2}b_2 \cosh(2\alpha_1\theta_0)e^{2\alpha_1^2\sigma^2T} \right). \end{aligned}$$

Let $T \rightarrow \infty$, $C^T - C^* \rightarrow 0$ by Assumption 1. ■