Dynamic Communication with Commitment

Yi Chen*

November 18, 2018

Abstract

I study the optimal communication problem in a dynamic principal-agent model. The agent observes the evolution of an imperfectly persistent state, and makes unverifiable reports of the state over time. The principal takes actions based solely on the agent's reports, with commitment to a dynamic contract in the absence of transfers. Interests are misaligned: while the agent always prefers higher levels of action to lower, the principal's ideal action is state-dependent.

In a one-shot interaction, the agent's information can never be utilized by the principal. In contrast, I show that communication can be effective in dynamic interactions, and I find a new channel, the *information sensitivity*, that makes dynamic communication effective. Moreover, I derive a closed-form solution for the optimal contract. I find that the optimal contract can display two properties new to the literature: contrarian allocation, and delayed response. I also provide a necessary and sufficient condition under which these properties arise. The results can be applied to practical problems such as capital budgeting between a headquarters and a division manager, or resource allocation between the central government and a local government.

Keywords: Communication; Dynamic contract; Contrarian; Delay; Prudence; Diffusion
JEL Codes: D82, D83, D86

*Cornell University, 106 East Ave, Ithaca, 14853. Email: yi.chen@cornell.edu. I am grateful to my dissertation committee: Johannes Hörner, Larry Samuelson, and Dirk Bergemann. I also thank Yeon-Koo Che, Chen Cheng, Eduardo Faingold, Simone Galperti, Yuan Gao, Marina Halac, Ryota Iijima, Navin Kartik, Nicolas Lambert, Elliot Lipnowski, Chiara Margaria, Dmitry Orlov, Gregory Pavlov, Ennio Stacchetti, Philipp Strack, Juuso Välimäki, Zhe Wang, Yiqing Xing, and John Zhu for helpful comments and suggestions. All 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) often has a motive to misguide the uninformed party (principal), 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 private 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?

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 pursuing 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, and the information about the state is totally wasted.

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 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

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

same continuation payoff, the principal has degrees of freedom to reallocate actions between the present and future in her favor. This is how a dynamic contract improves communication.

I discover a new channel, called the *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 prospects of the product. Then the slope of the target function, defined as information sensitivity, captures how much the ideal amount of resources marginally changes with the prospects. The higher the information sensitivity, the more the headquarters' ideal resource allocation depends on the manager's private information.

Suppose the prospects 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 increases in the prospects do not justify any additional resources, then the information sensitivity is low.

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 the prospects of the product, 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 inter-temporal 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 prospects, 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 prospects augurs higher future prospects 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 contract is *contrarian*: decrease the current action when the reported state is high, and increase it when low. At the same time, the contract may exhibit a *lagged* response to the change in reported state: if the current state is claimed to be high, the action does not increase immediately, but it will in the future.

For the optimal contract to display those two properties, a necessary and sufficient condition is that the information sensitivity in the future is in expectation higher than it is now. Intuitively, suppose the current state increases. Because of the persistence in the state process, future states also increase. Ideally the principal would like to increase both the current action and the future actions. However, this is not incentive compatible for the agent. In order to induce truth-telling, the principal must increase one and decrease the other. Which direction leads to a profitable trade-off? When information sensitivity is expected to be higher in the future, the future is more important in the inter-temporal trade-off. As a result, the principal sacrifices the current action in order to increase future action. This leads to the contrarian allocation and the delayed response. 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. On the other hand, if the future information sensitivity is smaller, then the trade-off favors the current action. In this case, the allocation of actions moves along with the agent's report, and there is no delay. What if the future is exactly as important as the present? Then we are in the knife-edge case where there is no profitable trade-off in either direction. In this case, the principal optimally ignores the agent's information even though she can use the information in an incentive-compatible way. This contrasts to the reason why communication fails in a one-shot interaction.

Exactly when is future information sensitivity more likely to be higher, resulting in a contrarian mechanism? I show that the state process and the target function jointly determine the conditions. In particular, there are two forces at work. The first force lies in the time trend in the state. When the drift of the state process is high, and the information sensitivity marginally changes much with the state (the slope of information sensitivity function with respect to the state is high), then the future information sensitivity is likely

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

to be high. Second, it also comes from the uncertainty of the state process. When the volatility of the state is high, and the information sensitivity is convex in the state, then due to Jensen's inequality, the future information sensitivity is likely to be high. The second force relates to the precautionary saving literature, where prudence matters, and marginal utility corresponds to the information sensitivity in this paper.

The contrarian property also relies on the persistence of the state process. If the state displays strong mean reversion, then the distribution of the future states depends little on the current state, and there is almost no reason to sacrifice current action in exchange for a better match in the future. This is why contrarian contract does not arise in the literature where the state process follows a two-state irreducible Markov chain. On one hand, with only two states, the target function is linear by definition. On the other hand, mean-reversion is automatically built in with the irreducible 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.

The optimal contract has a simple implementation. In the beginning, the principal assigns the agent a fixed total budget of actions, and commits to it. Each period, the agent reports the state to the principal, and the principal optimally chooses how much actions to assign. The more actions are used today, the less remains in the pool for future use. Commitment power kicks in only for keeping the total budget of actions fixed. Other than that, the principal' choice is sequentially rational.

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.

Related Literature 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 or dynamic cheap talk and 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 of allocation. They introduce a quota mechanism that takes advantage of the Law of Large Numbers to achieve asymptotic efficiency. 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. Margaria and Smolin (2017) prove a version of folk theorem in a repeated cheap talk game with multiple senders. 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. Koessler and Martimort (2012) study a static delegation problem with two-dimensional decision space, where valuable delegation arises as the principal uses the spread between the two decisions for screening purpose. 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. Malenko (2016) examines a dynamic capital budgeting problem with costly verification, and finds the optimal mechanism to be a inter-temporal budget with threshold separation of financing. In these papers, contrarian 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 inter-temporal trade-offs that shape a dynamic contract, it is convenient to start with a static example and then contrast it with a two-period example.

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

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

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

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

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

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

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

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

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

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

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

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

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

The optimal contract reads (derivation relegated to Appendix):

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

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

To interpret, both x_1^* and x_2^* can be decomposed into a term responsive to θ_1, and a constant term. The pair of actions add up to $W^* \equiv x_1^*(\theta_1) + x_2^*(\theta_1) = \mathbb{E}f(\theta_1) + \mathbb{E}f(\theta_2)$, meaning that the optimal choice of the total payoff is the unconditional expectation of total targets across periods[3]. Also, for either period, the constant term is simply half of W^*, a result from "cost smoothing".

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

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

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

If the first effect is stronger, then it is more cost-efficient to let the action x_1^* move in the *same* direction as the target $f(\theta_1)$, at the cost of worse expected match in period 2. The action is thus called *conformist* at state θ_1. If the second effect is stronger, then cost-efficiency requires the action x_1^* to move in the *opposite* direction as the target as a sacrifice, in exchange for a better match in period 2. If this is the case, then the action is called *conformist* at state θ_1. For instance, when $f(\theta) = e^\theta$, the action is contrarian at all θ_1. For target function $f(\theta) = \theta - \frac{1}{2}|\theta|$, contrarian action occurs only at a subset of states. The action is always conformist if $f(\theta) = \frac{|\theta|^{5/2}}{\theta}$. Moreover, if $f(\theta) = \theta$ or $f(\theta) = \theta^2$, then the two

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

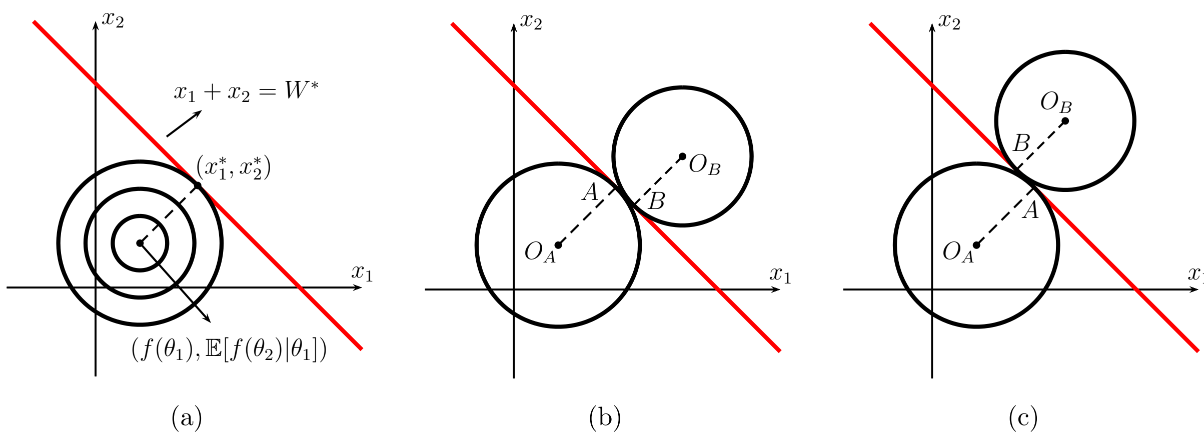

Figure 1: Iso-cost curves and the budget line. Panel (a): the family of iso-cost curves and the budget line. Panel (b): conformist—x_1 and $f(\theta_1)$ move in the same direction. Panel (c): contrarian—x_1 and $f(\theta_1)$ move in opposite directions.

terms in the bracket of (5) always cancel out. As a result, the action pair does not respond to information at all, indicating a failure of communication.

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

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

3 Continuous-Time Model: Settings

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

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

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

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

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

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

Interests are severely misaligned. While the principal's favorite action depends on the state, the agent only wishes to induce actions as high as possible. Specifically, the principal's flow cost at time t from a state-action pair (θ_t, x_t) is $(x_t - f(\theta_t))^2$, i.e., she suffers a quadratic cost from the gap between the actual action x_t and the state-dependent *target* action $f(\theta_t)$.[4] The agent's flow payoff is simply x_t, independent of the state.[5] The linear payoff of the agent shuts down risk aversion as a possible channel of rent extraction, as has been done in the literature.

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

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

whenever well-defined.

[4]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 form with re-scaling. In general cases, the quadratic cost is often a good approximation.

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

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

The principal commits to a contract at time zero. I use superscript t to denote a history up to time t. A contract x is a $\hat{\theta}$-measurable process specifying an action $x_t(\hat{\theta}^t) \in \mathbb{R}$ as a function of the report history, for all $t \geqslant 0$. There is no transfer of money.[6] A strategy of the agent is a θ-measurable process m. It prescribes the drift $m_t(\theta^t)$ that the agent adds to the true state process as a function of the state history, for all $t \geqslant 0$. I define the space of feasible strategies as:

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

to exclude explosive strategies. This restriction is not essential; In the end of Appendix I show the effect of relaxing the strategy set.

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

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

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

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

4 Incentives of the Agent

This section performs two standard steps of simplification, facilitating the analysis of the optimal contract in Section 5. First, I restrict attention to truthful contracts by invoking a version of the Revelation principle. Second, I use the first-order approach to derive a necessary condition for incentive compatibility. Sufficiency is verified after the candidate solution is obtained in Section 5.

[6]Unconstrained transfer of money trivializes the problem with the first-best solution. Money transfer with limited liability effectively serves to reshape the principal's loss function, which is discussed in Section 6.

4.1 Revelation Principle

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

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

Proof. See Appendix. ■

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

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

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

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

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

4.2 Incentive Compatibility: Necessary Condition

The first-order approach in the literature (Williams (2011), Kapička (2013), Pavan, Segal, and Toikka (2014), DeMarzo and Sannikov (2016)) derives a local version of the incentive constraints, namely, conditions to prevent the agent from local deviations.

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

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

as the agent's on-path expected continuation payoff. The expectation $\mathbb{E}_t$ is conditional on the information generated by the state by time t. As is verified later, this W_t together with

θ_t suffices to summarize the history up to t.[7] From now on, I suppress the dependence of W on x to simplify notations.

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

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

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

Proof. See Appendix. ■

The first term on the right-hand side of (8)represents the drift of W_t: it grows at rate r and depletes as x_t is paid out to the agent. The second term, which is the diffusion, governs the incentives. On equilibrium path, it holds that $\mathrm{d}\hat{\theta}_t - \mu\mathrm{d}t = \sigma\mathrm{d}Z_t$, which has zero mean. The multiplier $r\beta_t$ is interpreted as the instantaneous slope of the continuation payoff with respect to reported states, or "strength of incentives" (He, Wei, Yu, and Gao (2017), DeMarzo and Sannikov (2016)). Suppose the instantaneous slope $r\beta_t$ is positive. By adding a drift $m > 0$ to the report for a short moment $\mathrm{d}t$, the agent tricks the principal into believing that the state is $m\mathrm{d}t$ higher than it actually is, and this results in an increase of $r\beta_t m\mathrm{d}t$ in the agent's continuation payoff.[8] Similarly, if $r\beta_t < 0$, the agent can profit by shading the report. The only way to deter a local deviation from truth-telling is to keep $r\beta_t$ identically at zero. Proposition 1 formalizes the above reasoning as a necessary condition for incentive compatibility.

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

Proof. See Appendix. ■

According to the proposition, the only way to induce truth-telling from the agent is to entirely disentangle his continuation payoff from the current report, a theme already foreshadowed in the two-period model. While it may seem severe as a constraint for the

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

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

principal, there are lots of degrees of freedom to maneuver: a given continuation payoff can be supported by infinitely many paths of actions. Therefore, the choice of how the action paths responds to information, subject to the IC-FOC, is the key to optimal utilization of information, and that is the task of Section 5.

5 Optimal Contract

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

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

Assumption 1 (Regularity)
(i) The target f is piecewise $\mathcal{C}^2$;
(ii) There exists $\alpha_0 > 0, \alpha_1 \in \left[0, \frac{\sqrt{\mu^2+2r\sigma^2}-\mu}{2\sigma^2}\right)$ such that $|f(\theta)| \leqslant \alpha_0(e^{\alpha_1\theta} + e^{-\alpha_1\theta})$.

Part (i) of the assumption puts some smoothness on the target function to enable local analysis. Part (ii) uniformly bounds the target by some exponential function with a low growth rate. It ensures the target does not drift to infinity too fast, necessary for costs and payoffs to be finite.

5.1 Solving for the Optimal Contract

I use a recursive formulation to solve the relaxed problem where only the IC-FOC is considered. As is conjectured in Section 4.2, the optimal contract can be written in terms of two state variables: the state θ and the continuation payoff W. In Theorem 1 I formally prove that the candidate contract derived in the recursive problem is indeed the solution to the original problem (6)-(7).

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

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

The right-hand side of the equation consists of four terms. The first term is the normalized flow cost. The second and third terms are expected changes in cost caused by the drift in W and θ. The last term, the Itô term, reflects the volatility in θ. In order for the recursive

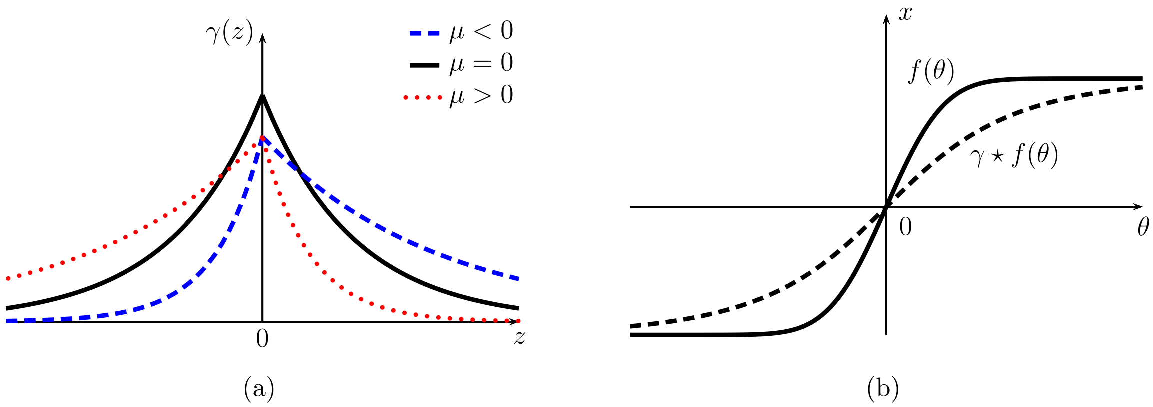

Figure 2: The kernel and the convolution. Panel (a): Graphs of the kernel γ for $\mu < 0$, $\mu = 0$, and $\mu > 0$. Panel (b): The target function f (solid curve) and the convolution $\gamma \star f$ (dashed curve) for $\mu = 0$.

form to make sense, the cost and the payoff must also satisfy the transversality condition:

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

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

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

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

where γ is a kernel with asymmetric Laplace distribution:

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

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

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

Theorem 1 (Optimal Contract)

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

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

Proof. See Appendix. ∎

Theorem 1 has several implications. First, since the principal is free to choose any initial W_0, the cost is minimized at $W_0 = \gamma \star f(\theta_0)$, which is the expected discounted future targets. That is, the principal should not distort the actions *on average* in the optimal contract; instead, the optimality comes from the way in which the action responds to the state. This will be discussed in detail in Subsection 5.2.

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

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

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

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

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

Theorem 2 (Impossibility)
(i) Zero cost is obtained if and only if the target is almost everywhere identical to a constant.
(ii) For $\mu = 0$, the cost reaches the babbling cost if and only if the target is almost everywhere

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

Proof. See Appendix. ■

The theorem is interpreted as two impossibility results. First, friction-less communication is not achievable unless the state has no bearing on the principal's cost. Second, for some non-generic set of target functions, communication fails even with a dynamic contract. Specifically, if the target is affine in the state, then the action is a constant over time and babbling is the inevitable outcome. Stranger still, babbling is optimal even when the target has some particular form of curvature. Theorem 2 conveys the message that it is necessary to investigate higher order derivatives of the target than simply the curvature. This is explored further in the next subsection.

5.2 Conformist or Contrarian Response?

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

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

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

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

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

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

Theorem 3 (Conditions for Conformist and Contrarian)
The following statements are equivalent:

(i) The optimal contract stipulates a conformist (contrarian, resp.) action at state θ.
(ii) $(f'(\theta) - (\gamma \star f)'(\theta))f'(\theta) > 0$ $(< 0,\ \textit{resp.})$.
(iii) $\Big((\gamma \star f)''(\theta), (\gamma \star f)'''(\theta)\Big) \cdot (2\mu, \sigma^2)f'(\theta) < 0$ $(> 0,\ \textit{resp.})$.

Proof. See Appendix. ■

According to the theorem, contrarian actions can occur in many scenarios. Subsection 5.4 is dedicated to enumerating economically meaningful examples where contrarian actions arise. Why would contrarian action ever be optimal? After all, it is tempting to increase the action following a rise in the target, the usual way that the "quota" of actions is used in the literature.

Here is the intuition. Heuristically, the slope $f'(\theta)$ can be interpreted as the *information sensitivity* of the principal at state θ. A steeper slope corresponds to a target that is more sensitive to a state change. Suppose $f' > 0$ for simplicity. If the current state has a shock $\mathrm{d}\theta > 0$, then the current target increases by $f'(\theta)\mathrm{d}\theta$,

and the principal is tempted to increase the current action. Meanwhile, because of the persistence of the state process, the expected discounted future target increases as well, by $(\gamma \star f)'(\theta)\mathrm{d}\theta$. This also creates the motive to increase total future actions to better match the higher expected targets. However, she cannot achieve both due to incentive constraints: higher current action must be followed with lower future actions, and vice versa. Therefore, if $f'(\theta)$ is larger, then the current information sensitivity is higher than future, and the inter-temporal trade-off tips the scale in favor of the current action. In this way, the current action increases along with the current target (conformist), at the cost of future matching. Conversely, if the future information sensitivity $(\gamma \star f)'(\theta)$ is higher, then the principal optimally sacrifices the current matching in order to rebalance the Euler equation. As a result, the current action moves against the current target (contrarian).

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

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

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

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

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

5.3 Evolution of the Contract

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

Proposition 2 (Cost and Payoff Dynamics)
(i) The continuation cost is a sub-martingale, i.e., $\frac{\mathbb{E}_t\left[\mathrm{d}C_t\right]}{\mathrm{d}t} \geqslant 0$*;*
(ii) The continuation payoff monotonically increases (decreases, resp.) over time if $\gamma\star f(\theta) - f(\theta) > 0$ *(*< 0*, resp.) for all* θ*.*

Proof. See Appendix. ■

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

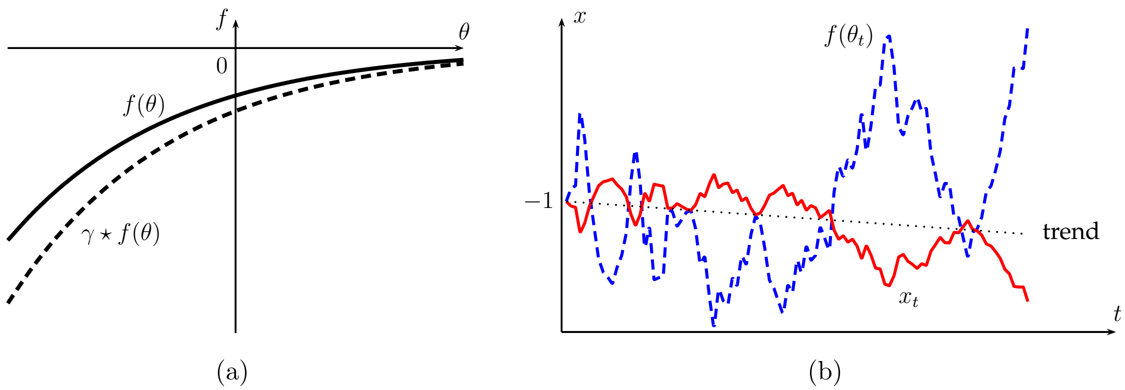

Figure 3: Exponential target with parameters: $r = 1, \sigma = 1, \mu = 0, \theta_0 = 0, b_0 = -1, b_1 = -0.7$. Panel (a): the target function (solid curve) and its convoluted version (dashed curve). Panel (b): Simulated paths of the target (blue dashed curve) and the optimal action (red solid curve). The dotted curve is the same trend for both paths; it always separates the target and the action on opposite sides.

distortion accumulates over time on average. It is shown in the proof that the continuation cost C_t always has a non-negative drift, strictly positive if the target is not a constant.

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

5.4 Examples

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

5.4.1 Exponential Target

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

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

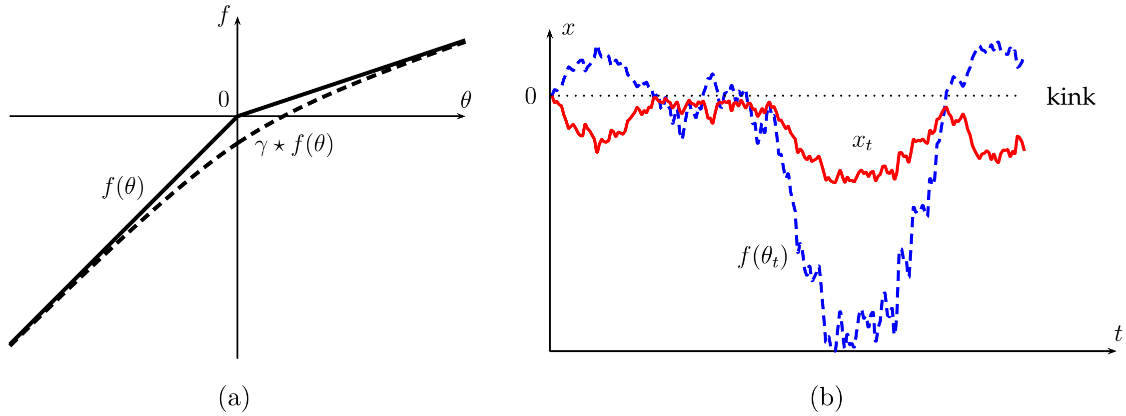

Figure 4: Kinked target with parameters: $r = 1, \sigma = 1, \mu = 0, \theta_0 = 0, b = 0.33$. Panel (a): the target function (solid curve) and its convoluted version (dashed curve). Panel (b): Simulated paths of the target (blue dashed curve) and the optimal action (red solid curve). The dotted curve shows the locus of the target at the kink. The cyclicality of the two paths depends on whether the target is above the kink.

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

5.4.2 Kinked Target

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

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

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

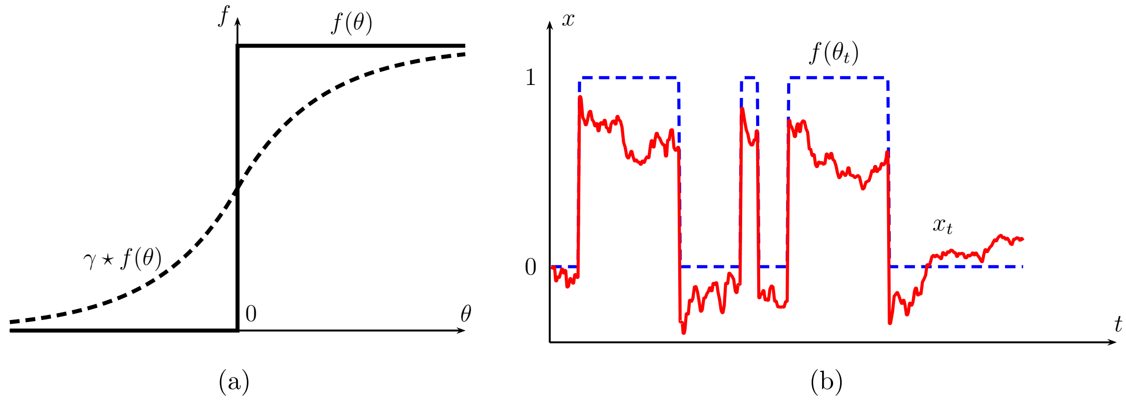

Figure 5: Binary target with parameters: $r = 1, \sigma = 1, \mu = 0, \theta_0 = -0.4$. Panel (a): the target function (solid curve) and its convoluted version (dashed curve). Panel (b): Simulated paths of the target (blue dashed curve) and the optimal action (red solid curve). The action always jumps one-for-one when the target jumps, otherwise it always moves in the opposite direction of the state.

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

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

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

5.4.3 Binary Target

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

To be specific, let $f(\theta) = \mathbb{1}\{\theta \geqslant 0\}$. While the target is discontinuous at 0, the expected future target $\gamma \star f$ is continuously differentiable. It can be shown that $f' < (\gamma \star f)'$ at all states except zero, the discontinuity point. Since the target is flat at these states, the action is neither conformist nor contrarian. However, the action *does* respond to the state in the opposite direction. Only at the cutoff state 0 does the action jump in the same direction and with the same magnitude as the target. Panel (a) of Figure 5 pictures the target and the expected future target. Panel (b) simulates the time paths for the target and the action. Every time the target jumps, the optimal action jumps one-for-one. This happens when the state crosses the cutoff 0. At other times, the action moves against the state, which is not shown in the picture.

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

6 Extensions

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

6.1 Mean Reverting State Process

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

When the state process exhibits mean-reversion, the persistence is weaker. For simplicity, here I assume a mean-reverting state process:

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

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

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

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

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

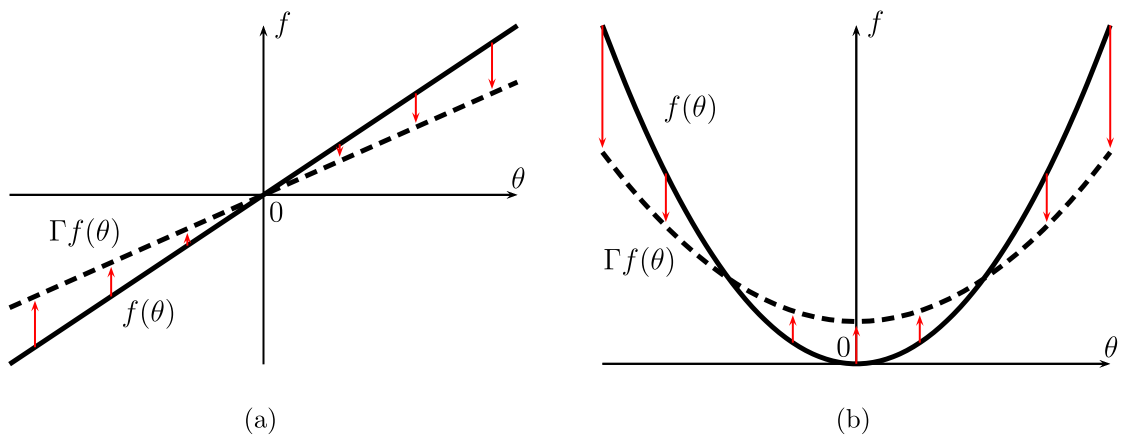

Figure 6: The effect of mean-reversion, with parameters $r = 1, \sigma = 1, \phi = .5, \theta_0 = 0$. Panel (a): The solid curve is the target $f(\theta) = \theta$ and the dashed curve is Γf. Panel (b): The solid curve is the target $f(\theta) = \theta^2$ and the dashed curve is Γf.

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

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

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

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

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

6.2 Transfer with Limited Liability

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

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

$$\begin{aligned}
\min_{\substack{(x_t(\cdot))_{t\geqslant 0} \\ (y_t(\cdot)\geqslant 0)_{t\geqslant 0}}} \quad & \mathbb{E}\left[\int_0^\infty re^{-rt}\left((x_t(\theta^t) - f(\theta_t))^2 + y_t(\theta^t)\right)\mathrm{d}t\right] \\
\text{s.t.} \quad & \mathbb{E}\left[\int_0^\infty re^{-rt}(x_t(\theta^t) + y_t(\theta^t))\mathrm{d}t\right] \geqslant \mathbb{E}\left[\int_0^\infty re^{-rt}(x_t(\hat{\theta}^t) + y_t(\hat{\theta}^t))\mathrm{d}t\right], \quad (15) \\
& \text{where } \hat{\theta}_t = \theta_t + \int_0^t m_s \mathrm{d}s, \ \forall\, m \in \mathcal{M}.
\end{aligned}$$

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

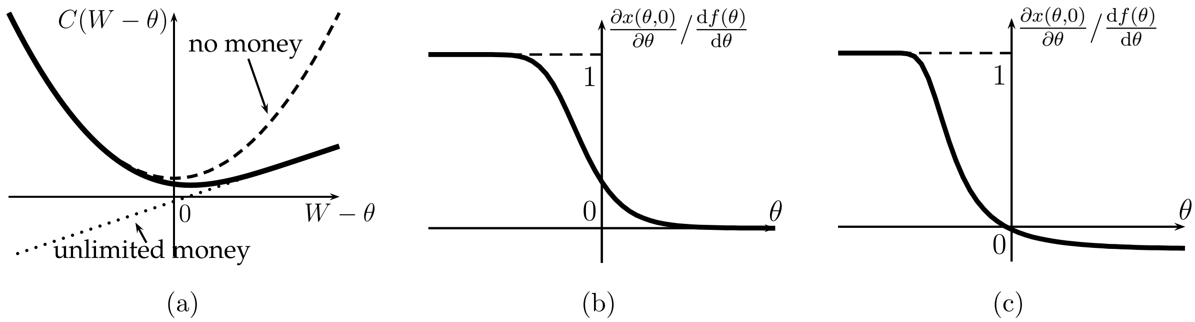

Figure 7: Transfer with limited liability. Parameters: $r = 1, \sigma = 1$. Panel (a): cost function in solid curve, with $f(\theta) = \theta$. Panel (b): the ratio of derivatives at $W = 0$ in solid curve, with $f(\theta) = \theta$. Panel (c): the ratio of derivatives at $W = 0$ in solid curve, with $f(\theta) = -e^{-0.6\theta}$.

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

Similar results extend to nonlinear target functions. Panel (c) of Figure 7 shows the ratio of derivatives holding $W = 0$, for the target function $f(\theta) = -e^{-0.6\theta}$. Again, its behavior approaches either contrarian or perfect conformist based on the tightness of money.

7 Conclusion

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

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

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

References

Abreu, 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. Sonnenschein (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.

Koessler, F., and D. Martimort (2012): "Optimal delegation with multi-dimensional decisions," *Journal of Economic Theory*, 147(5), 1850 – 1881.

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 multistage communication," *Journal of Economic Theory*, 117(2), 147 – 179.

Malenko, A. (2016): "Optimal dynamic capital budgeting," *Working paper*.

Margaria, C., and A. Smolin (2017): "Dynamic communication with biased senders," *Games and Economic Behavior*.

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

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.