Relational Influence

Heikki Rantakari*
University of Rochester
Simon Business School

October 11, 2016

Abstract

An uninformed principal elicits non-contractible recommendations from a privately informed agent regarding the quality of projects. The agent is biased in favor of implementation and no credible communication is possible in a one-shot setting. In a repeated setting, the fear of losing future influence can sustain informative communication, but the agent's willingness to remain truthful depends on the extent to which he expects the principal to listen to him. In a stationary equilibrium, the principal always implements mediocre projects at a sub-optimally high frequency to reward honesty, while she may either favor or discriminate against high-quality projects. In a non-stationary equilibrium, the principal will further condition the agent's future influence on today's proposals, with the admission of mediocre alternatives rewarded with increased future influence while rejections of high-quality proposals are further punished by lowering the agent's future influence.

*contact: heikki.rantakari@simon.rochester.edu. I would like to thank Michael Raith and Luis Rayo for helpful comments. The usual disclaimer applies.

1 Introduction

Issues of decision-making under strategic information transmission have been increasingly recognized to be of crucial importance for organizational performance. As noted by Cyert and March (1963 [1992]), "[w]here different parts of the organization have responsibility for different pieces of information..., [we would expect] attempts to manipulate information as a device for manipulating the decision." (p.79). Following the theoretical frameworks introduced by Crawford and Sobel (1982) and Milgrom and Roberts (1988), a large and growing literature has examined how incentive conflicts lead to attempts at manipulating information, the resulting loss of information and how to manage such losses, leading to the analysis of issues of delegation, mediation, information management and beyond.

A particular aspect of such relationships, especially in organizational settings, is that they are typically ongoing. For example, a company does not just choose a single R&D project, build a single factory or choose where to locate a new logistics center. Instead, the organization faces an ongoing sequence of such decisions, and it is the same group of organizational members that are involved in the decision-making process. This ongoing nature of the decision-making process then allows the development of relationships, and even if the quality of current recommendations cannot be verified on the spot, the parties can learn about the quality of past recommendations from the outcomes that have resulted.

This paper constructs a simple model of such relationships and considers how the parties can use the history of the relationship as a basis for current behavior and to sustain a relationship that is better for both parties than a one-shot interaction. In the setting, a principal needs to decide whether to implement a project. The project can be either mediocre or good, information learned only by the agent. The agent makes a recommendation to the principal regarding the quality of the project, after which the cost of implementation is publicly observed. The principal wants to implement the project only when its value exceeds the cost of implementation, while the agent is biased in favor of implementation. If the principal chooses to implement the project, the value of the project is learned before the next choice is made, while if the principal chooses against implementation, the value of the project is not learned. This asymmetry captures the idea that we learn less about recommendations that are not followed relative to recommendations that are followed.

The object of interest is the principal's decision rule, which determines when the project is implemented, conditional on whether the agent submits either a weak or a strong report and the realized cost of implementation.[1] Since this decision rule effectively determines how likely it is that the agent's proposal is implemented, it determines the agent's influence in the relationship. The goal of the analysis is to examine how the principal can use this relational influence of the agent to manage the relationship.

The first observation is that trust creates value because information is valuable to both parties. When the agent is sufficiently patient, he is willing to truthfully reveal the quality of the project

[1] To focus the analysis on decision-making, I assume that direct monetary transfers are not available. The results are analogous if we allow only for one-sided transfers from the principal to the agent. The agency problem disappears under deep pockets.

and the principal can follow the first-best decision rule. This result follows because if the agent misleads the principal, the principal will learn that after the fact and will stop trusting the agent's recommendations going forward, leading to less informed decision-making and worse payoffs to both parties. But when the agent is not patient enough, the temptation to push for the acceptance of a mediocre project becomes too high and the first-best outcome is no longer implementable. The question then becomes what the principal can do to maintain the relationship.

Despite the simplicity of the basic setting, the equilibrium displays a rich level of interactions. Thus, I begin by considering a stationary equilibrium, where the current decision rule is not a function of past history but can contain biases both in favor and against a particular proposal. Noting that the main constraint that we need to satisfy is keeping the agent honest about mediocre projects, the basic distortions are two-fold. First, the principal will always bias the acceptance rule in favor of mediocre alternatives and implements some mediocre projects at a loss. The reason for this result is simple: by implementing the mediocre project more frequently, the agent gains less from exaggeration and thus makes him less willing to sacrifice his future influence for an immediate gain. Second, the principal will either favor or discriminate against high-quality projects. This result is the effect of two counteracting forces. On one hand, favoring high-quality projects increases the agent's future influence and thus makes the agent less willing to sacrifice that influence for an immediate gain. On the other hand, favoring high-quality projects also increases the gains from exaggeration by making the strong proposal more influential relative to a weak proposal. When the agent is sufficiently patient, the first effect (future value) dominates, while when the agent is less patient, the second effect (higher immediate gain) dominates. Thus, an agent of intermediate patience is rewarded with higher than the first-best level of influence (with the overall decision rule biased in favor of implementation), while an agent with lower patience has lower than the first-best level of influence – while the agent gets mediocre proposals implemented more frequently, the fact that he is discriminated against when his project is of high quality more than offsets that gain.

Finally, while both types of projects may be implemented with excessive frequency, the relative likelihood that the low-quality project is implemented is monotone increasing in the impatience of the agent. In practical terms, the stationary equilibrium thus provides a simple explanation for corporate socialism, whereby the internal allocation of resources is less responsive to differences in profitability as suggested by simple NPV criterion. This lack of responsiveness is used by the principal to satisfy the truth-telling constraint of the agent. Relatedly, the variation in the patience of the agent leads to either over-investment or under-investment relative to the first-best allocation.

Having considered the stationary equilibrium, I then consider how the principal can do better by using a non-stationary decision rule and altering the agent's future influence based on today's recommendation and the resulting decision. The basic results are three-fold. First, the principal will reward the agent for admitting that his project is mediocre by increasing his future influence. The benefits of this reward are two-fold. First, because the cost of distorting the decision is convex in the distortion, spreading the reward for honesty over several periods lowers the overall cost of the reward. Second, because the agent also values the implementation of high-quality projects more, a promise of future influence is a more efficient means of rewarding the agent – instead of settling up now by implementing a mediocre project, the principal promises to give the agent more favorable treatment in the future, when the agent may have a high-quality project available. This result thus provides a simple rationale for basic quid pro quo arrangements, where honesty today is rewarded by favorable treatment in the future. For example, a department admitting that their favorite job candidate is mediocre is promised priority in the hiring process next year, or a division admitting for limited investment opportunities today is promised easier access to funding for any new projects the following year.

Second, in addition to rewarding the admission of mediocre alternatives, the principal will also punish the recommendation of high-quality alternatives by lowering the agent's future influence if the recommendation is rejected. The reason for this result is that when a proposal is not accepted, its quality is not learned. Thus, a strong recommendation is not tested if it is rejected, which limits the punishment available in the case of deviation. By lowering the agent's payoff when the proposal is rejected, the incentives to exaggerate are decreased. However, because even honest recommendations of high-quality projects are sometimes rejected (and, indeed, in equilibrium, only honest recommendations arise), terminating the relationship is too harsh of a punishment. Instead, the principal responds by lowering the future influence of the agent without fully stopping trusting him. In pratice, this feature resembles a situation where an agent falls out of favor with the principal – once a strong proposal is rejected, the agent's odds of getting future projects through are lowered.

Third, because of the convexity of the losses faced by the principal in the size of the distortions, the evolution of influence will be gradual. Admissions of mediocre projects will slowly increase influence while the rejections of strong proposals will gradually erode it. And because of the ongoing value of motivating the agent through changes in the continuation value of the relationship, there is no absorbing state to the game. Instead, the influence will fluctuate back and forth depending on the history of the play. Because of the complex interplay across the distortions in the decision rule of the stage game and the continuation values, solving for the optimal equilibrium is challenging. As a result, to illustate these basic results, I provide a simplified example of a three-stage equilibrium, where the stages can be ranked in terms of the level of influence by the agent and the transitions and the stage-game distortions are optimally managed by the principal to maximize the value of the relationship to her.

The challenge created by the distortions in the stage game, whether in the stationary or non-stationary equilibrium, is that as the agent becomes increasingly impatient, the distortions needed to keep the agent truthful grow. Then, the need for the principal to honor the promises made in equilibrium can put an upper cap on the distortions that can be sustained and the relationship may become unsustainable – the distortions needed to keep the agent truthful are too large to be credible for the principal to follow, and the whole relationship collapses.

The rest of the manuscript is organized as follows. Section 2 discusses the related literature and section 3 outlines the model. Section 4 provides a preliminary examination of the framework, illustrating the full set of payoffs attainable in the stage game and the conditions under which the first-best equilibrium can be obtained. Section 5 derives the optimal stationary equilibrium, and Section 6 considers dynamics. Section 7 concludes and discusses some potential extensions. Appendix B illustrates a game with a continuous project space for the agent to illustrate how the stationary decision rule is optimally determined in such settings.

2 Related Literature

This paper lies in the intersection of the literatures on repeated games and strategic communication. The five main papers existing in this intersection are Alonso and Matouschek (2008), Kolotilin and Li (2015), Campbell (2015), Li et al. (2015) and Lipnowski and Ramos (2015). The first two papers consider the classic repeated-game setting of full ex post observability of outcomes. Alonso and Matouschek (2008) consider the Crawford and Sobel (1982) setting with a long-lived principal interacting with a sequence of myopic agents. The value of the ongoing relationship helps the principal to choose a decision closer to the agent's preferences, facilitating communication. If the principal is sufficiently patient, she achieves the optimal delegation set (and thus the maximum payoff to the principal in the absence of transfers). Kolotilin and Li (2015) extend the CS setting to a long-lived agent and the ability to make transers. Transfers make the communication problem trivial by having the agent signal his information with an associated transfer, and the issue is how to manage decision-making by the principal who underweighs the agent's payoff in her favored decision. I consider a qualitatively different decision problem, which introduces the asymmetric learning regarding the quality of recommendations based on whether they are followed or not.

The remaining three papers consider settings that are qualitatively closer to the present model, with the agent making recommendations regarding which (if any) projects to implement, but consider the opposite extreme of no learning of the outcomes. The strategies can thus be based only on the observed history of recommendations. In Campbell (2015), the agent uses his relational capital to recommend a project as long as it is good enough, which uses his relational capital until it is exhausted. The capital is never replenished. Lipnowksi and Ramos (2015) is closest to the present paper, where the relational capital is both replenished and used over time, but where the replenishment occurs when the agent recommends rejection while the capital is used whenever the agent recommends acceptance.[2] Finally, in Li et al. (2015) the agent can recommend either his ideal project or a project that is better for the principal, but that project may not be available. When the agent recommends his own project, the continuation value must punish the agent to keep him honest, drifting the equilibrium towards not listening to him, while recommending the project that is better for the principal increases the continuation value, with an increased likelihood that the agent gets to choose his preferred project whenever he wants to. However, in all three papers the question is simply whether the principal follows the agent's recommendation, thus not allowing for the richer manipulation of the acceptance rule to manage the relationship, which is the focus here.

More broadly, the present paper relates to the large literature on repeated games with private information that has followed Abreu et al. (1990), where the focus on the use of continuation values and distortions in the behavior (instead of monetary transfers) to sustain the equilibrium is present in, e.g. Athey and Bagwell (2001) and Athey et al. (2004) on colluding with private information, Hauser and Hopenhayn (2008) on favor trading, Li and Matouschek (2008) on enforcing the payment of bonuses by the principal, Padro i Miquel and Yared (2012) on managing the moral hazard problem of an intermediary in maintaining the rule of law, and Andrews and Barron (2014) on managing

[2] For a related paper, see Guo and Horner (2015).

multiple supply relationships, just to mention a few.

In terms of strategic communication, the paper considers a variant of the framework analyzed in Li et al. (2016), Rantakari (2016), Garfagnini et al. (2014) and Chakraborty and Yilmaz (2013), among others, where a decision-maker needs to choose among discrete alternatives, based on the recommendation(s) of an agent or multiple agents. The setting retains the discrete nature of the final choice, but introduces the continuity of private information for the principal.

3 Model

I consider a repeated advisory relationship between an agent and a principal. In the stage game, the agent has access to a "project," the value of which is given by $\theta \in \{\theta_L, \theta_H\}$, with $0 \leq \theta_L < \theta_H \leq 1$. Let the probability of the high-quality alternative be given by p. The agent observes privately the value of the project, and makes a recommendation $m_i \in \{m_L, m_H\}$ to the principal as to the quality of the project (since the project quality is binary, we can restrict our attention to binary messages). The recommendations are soft information (cheap talk), and the principal interprets the message according to equilibrium play to form beliefs regarding the quality of the project, $E(\theta|m_i)$.

Following the recommendation, the principal's outside option, c, is drawn from a known distribution F, which, for tractability, I assume to be $U[0,1]$. Once the outside option is realized and publicly observed, the principal chooses whether to adopt the project of the agent or choose the outside option. For concreteness, we can take the value to be the expected revenue generated by a given project, while the outside option is the cost of investment. Then, if the principal accepts the agent's project, the payoff is given by $\theta_i - c$ while the outside option is no investment, with normalized payoff of 0. The agent, on the other hand, does not bear any of the cost of investment, and the agent's payoff is given by θ_i if the principal invests and 0 if the principal doesn't invest. The principal observes her payoff at the end of the period, so she will learn whether the agent told the truth if she follows the recommendation, but does not if she chooses the outside option. The discount rates are $\delta_A, \delta_P < 1$ for the agent and the principal, respectively.

No transfers: If the parties had access to (unbounded) transfers, the solution would be trivial even in the static setting, as the agent could signal his private information through voluntary transfers. In many settings, however, transfers are either not available or are limited for various reasons, including risks of collusion and rent-seeking activities. Thus, I make the opposite assumption, where no transfers are available between the agent and the principal. Instead, the relationship will be sustained by the principal's decision rule, $\Pr(A|m_i, c)$, which specifies the (A)cceptance probability following the agent's message and the commonly observed principal's state. As noted by Cyert and March (1963), "Side payments, far from being the incidental distribution of a fixed, transferable booty, represent the central process of goal specification. That is, a significant number of these payments are in the form of policy commitments." (p.35)

Other assumptions: To maintain the tractability of the analysis and to be able to explore the dynamics of the relationship, I make a number of further simplifying assumptions. To mention a few, I assume a binary state for the private signal, publicly observable cost of investment, single agent, and perfect observability of the outcome when a project is implemented. The qualitative logic of the analysis remains if the agent's state is continuous, but the decision rule becomes naturally richer, highlighting differences among low- and medium-quality projects. This setting is discussed in Appendix B. The publicly observable principal's state allows us to focus on managing the relationship with only that one agent. An interesting avenue for future work is the examination of how to manage the relationship when multiple agents hold relevant information to the decision. Finally, a valuable extension would be to consider the imperfect observability of the outcomes, either by introducing noise into the principal's payoff or by allowing for imperfect information for the agent. However, the set of assumptions allows us to focus on how the principal can both instantaneously and dynamically manage the relationship with the agent when the only tool available is the decision rule, $\Pr(A|m_i, c)$.

4 Preliminaries - Feasible payoffs and First-Best

Before considering the equilibrium of the model, I will first consider the stage-game payoffs and the basic tradeoffs involved. This will help to summarize the structure of the model and thus provide insight into the results that follow later.

4.1 Payoff structure and the Pareto frontier

Begin by considering the expected payoff of the principal and the agent. Assuming truth-telling by the agent, we can write their payoffs as

$$u_P = \sum_{i\in\{L,H\}} \Pr(\theta_i)\Pr(A|\theta_i)\left(\theta_i - E\left(c|A,\theta_i\right)\right) \qquad \textit{and} \qquad u_A = \sum_{i\in\{L,H\}} \Pr(\theta_i)\Pr(A|\theta_i)\theta_i,$$

where $\Pr(A|.)$ indicates the probability of acceptance and $E\left(c|A,\theta_i\right)$ the expected cost of implementation, given the decision rule $\Pr(A|m_i, c)$ and the realized cost. As the first preliminary observation, note that since the agent cares only about implementation, the optimal acceptance rule will take a threshold structure, where the principal accepts the project as long as her cost of implementation is below some threshold, $\underline{c}(\theta_i)$. Given the assumption that the costs are uniformly distributed, we can then write the expected payoffs as

$$u_P = \sum_{i\in\{L,H\}} \Pr(\theta_i)\underline{c}\left(\theta_i\right)\left(\theta_i - \tfrac{\underline{c}(\theta_i)}{2}\right) \qquad \textit{and} \qquad u_A = \sum_{i\in\{L,H\}} \Pr(\theta_i)\underline{c}\left(\theta_i\right)\theta_i.$$

Second, note that the first-best threshold for implementation is $\underline{c}^{FB}\left(\theta_i\right) = \theta_i$. Thus, we can write any implementation rule simply as $\underline{c}\left(\theta_i\right) = \theta_i + x_i$, where x_i is the distortion away from the first-best decision rule. The principal's strategy can thus be summarized by the pair $\{x_L, x_H\}$, the examination of which will be the focus of the analysis. If, on the other hand, the principal makes no use of the available information, then the expected quality of the project is $E\left(\theta\right) = p\theta_H + \left(1-p\right)\theta_L$ and

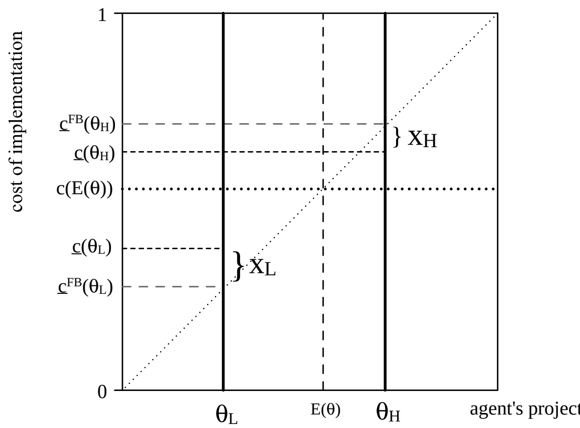

Figure 1: Illustrating the principal's decision rule

the optimal (common) threshold for implementation is then $\underline{c}^{FB}(E(\theta)) = E(\theta)$. In this case, the expected payoffs become simply

$$\underline{u}_P = \frac{E(\theta)^2}{2} \qquad and \qquad \underline{u}_A = E(\theta)^2.$$

The basic structure is illustrated in Figure 1. In this illustration, the principal is biasing her decision rule in favor of the agent when the agent is making the weaker recommendation, while discriminating against the agent when he makes the strong recommendation ($x_L > 0, x_H < 0$).

To measure the value of the agent's information, we can solve for $(u_P - \underline{u}_P)$ and $(u_A - \underline{u}_A)$, which are then given by

$$(u_P - \underline{u}_P) = \frac{1}{2}\left[(1-p)\,p\,(\theta_H - \theta_L)^2 - p x_H^2 - (1-p)\,x_L^2\right] \tag{1}$$

$$(u_A - \underline{u}_A) = p\,(1-p)\,(\theta_H - \theta_L)^2 + p\theta_H x_H + (1-p)\,\theta_L x_L \tag{2}$$

The value of the agent's information to the principal (and to himself) is thus proportional to $(1-p)\,p\,(\theta_H - \theta_L)^2$. The value is lowered for the principal, however, whenever the decision rule is distorted away from the first-best, so that $x_i \neq 0$. The cost of distortions is naturally convex in the size of the distortions, with the loss given by $p x_H^2 + (1-p)\,x_L^2$. The agent, on the other hand, benefits from a more favorable decision rule, where the value to the agent is given by $p\theta_H x_H + (1-p)\,\theta_L x_L$. I will call this component the relational influence of the agent, as any increase in x_i makes the principal more likely to follow the suggestion of the agent and thus increase his payoff.

The second question is what is the overall set of feasible payoffs. To this end, we must solve for the Pareto frontier, consisting of maximizing $(u_P - \underline{u}_P)$ conditional on delivering a given value

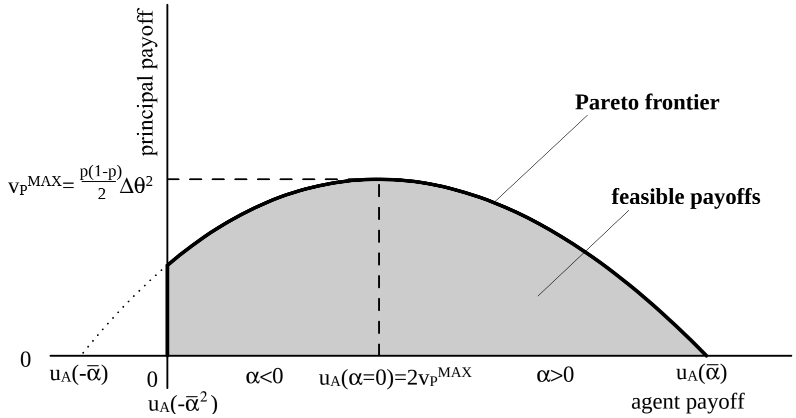

Figure 2: Pareto frontier and feasible payoffs

$(u_A - \underline{u}_A)$ to the agent. This involves solving for the least-cost deviation needed to deliver a given value to the agent. From equations 1 and 2 we obtain the following Lemma:

Lemma 1 ***Optimal distortions:*** *distortions* (x_L, x_H) *are on the Pareto Frontier if and only if* $\frac{\theta_L}{\theta_H} = \frac{x_L}{x_H}$. *We can thus characterize the Pareto frontier with a single coefficient* α, *where* $x_i = \alpha\theta_i$.

Proof. Holding the agent's expected payoff constant gives the tradeoff between the distortions as $\frac{dx_H}{dx_L} = -\frac{(1-p)\theta_L}{p\theta_H}$ while maximizing the principal's payoff requires $px_H \frac{dx_H}{dx_L} + (1-p)\, x_L = 0$. Together, these give $\frac{\theta_L}{\theta_H} = \frac{x_L}{x_H}$, which means any efficient distortion must satisfy $x_i = \alpha\theta_i$. ■

An important converse of this lemma is that if the decision rule is not proportional, then we are bounded away from the Pareto frontier. To finish characterizing the frontier, define $\overline{\alpha} = \sqrt{\frac{(1-p)p(\theta_H - \theta_L)^2}{p\theta_H^2 + (1-p)\theta_L^2}} < 1$ as the maximum level of influence that can be given to the agent in any game, defined by $(v_P(\overline{\alpha}) - \underline{v}_P) = 0$. Finally, note that (i) $u_P(\alpha = 0)$ gives the principal's preferred equilibrium, (ii) $\frac{dv_P(\alpha)}{dv_A(\alpha)} = \alpha$, so that the frontier is concave, and that (iii) each player is guaranteed a net payoff at least as high as zero. The resulting frontier together with all feasible payoffs is illustrated in Figure 2. The frontier is a quadratic equation with a unique maximum at the non-distorted decision rule, with the shaded area giving all feasible payoffs attainable in the stage game with suboptimal combinations of x_H and x_L.

4.2 Obtaining the first-best (and other parts of the Pareto frontier)

Consider now whether obtaining the principal's preferred equilibrium (first-best) is feasible. Since the maximizer for the principal is unique (at $\alpha = 0$), it can be obtained only if it is self-generating. Given that this is the principal's preferred equilibrium, we do not need to worry about her incentive-compatibility constraint. The only relevant constraint is for the agent to make the correct recommendation, which in turn can be binding only when faced with a mediocre alternative.

By telling the truth, the agent guarantees himself an expected payoff of

$$V_A = \theta_L^2 + \delta_A V_A. \tag{3}$$

In other words, his proposal is accepted with probability θ_L, and whether or not the proposal is accepted, the principal continues to trust the agent, keeping the game on the same continuation path, with value V_A. In contrast, if he chooses to deviate, his expected payoff is given by

$$V_A' = \theta_H \theta_L + \delta_A \left[\theta_H V_A^{dev} + (1 - \theta_H) V_A\right]. \tag{4}$$

He thus increases the immediate acceptance probability to θ_H by exaggerating the quality of the proposal, but now the lie is detected with probability θ_H, in which case the principal stops trusting the agent. However, if the proposal is still rejected, the principal learns nothing and thus the game remains on the equilibrium path, with value V_A. The strategies thus constitute an equilibrium if and only if

$$\frac{(\theta_H - \theta_L)}{\theta_H} \theta_L \leq \delta_A \left[V_A^{dev} - V_A\right], \tag{5}$$

which, in the case of $\alpha = 0$, simplifies to

$$\frac{\theta_L}{\theta_L + p(1-p)\theta_H(\theta_H - \theta_L)} \leq \delta_A. \tag{6}$$

Thus, as long as the agent is sufficiently patient, the principal is able to obtain her preferred outcome, simply by utilizing her preferred decision rule and stopping to trust the agent whenever he is caught lying. The higher the value of the relationship, $p(1-p)(\theta_H - \theta_L)^2$, the less patient the agent needs to be to remain truthful.

The more interesting case arises when the condition is not satisfied, and so the first-best outcome is no longer obtainable by the principal. The broader analysis is undertaken below, and I will only make two preliminary observations regarding the Pareto frontier. First, if the first-best is not obtainable at a given (δ_A, δ_P), then no point with $\alpha < 0$ is self-generating by itself since it provides a lower payoff to both the agent and the principal (however, it may be obtainable in some other equilibrium). Second, while the first-best may not be obtainable at a given (δ_A, δ_P), points with

$\alpha > 0$ may be self-generating. In particular, let $\phi(\theta) = (1-p)\,p\,(\theta_H - \theta_L)^2$ denote the value generated by information and $\eta(\theta) = \left(p\theta_H^2 + (1-p)\,\theta_L^2\right)$ the rate of value transfer as we alter α, then a point on the frontier will be self-generating as long as

$$\delta_P \geq \frac{2\,|\alpha|\,\theta_H}{2\,|\alpha|\,\theta_H + [\phi(\theta) - \alpha^2\eta(\theta)]} \qquad \textit{and} \qquad \delta_A \geq \frac{(\theta_H - \theta_L)\,\theta_L}{(\theta_H - \theta_L)\,\theta_L + \theta_H\,[\phi(\theta) + \alpha\eta(\theta)]}. \tag{7}$$

Note that as we increase α, while the equilibrium becomes more sustainable for the agent, it becomes less likely to be incentive-compatible for the principal. Once both parties are sufficiently impatient, no point on the frontier is self-generating.

5 Stationary equilibrium

Having considered the basic structure of the problem, we can now consider the repeated game itself. In this section, I will consider the optimal stationary equilibrium for the principal, where the distortions (x_L, x_H) are independent of the history of the play. In the next section, I will consider how the principal can do better by considering history-dependent strategies and what are the basic tradeoffs involved.

The distortions need to satisfy two incentive-compatibility constraints. First, as above, truth-telling needs to be in the agent's self-interest. To this end, let $u_A(x_L, x_H)$ denote the stage-game payoff for the agent and $\underline{u}_A$ the agent's payoff when caught misleading the principal. Then, we can write the agent's truth-telling constraint as

$$\begin{gathered} \Pr(A|m_L)\theta_L + \tfrac{\delta_A}{1-\delta_A} u_A(x_L, x_H) \geq \\ \Pr(A|m_H)\theta_L + \tfrac{\delta_A}{1-\delta_A}\left[\Pr(A|m_H)\underline{u}_A + (1 - \Pr(A|m_H))\,u_A(x_L, x_H)\right] \\ \Leftrightarrow \end{gathered}$$

$$\frac{\delta_A}{1-\delta_A}\left[u_A(x_L, x_H) - \underline{u}_A\right] \geq \frac{(\Pr(A|m_H) - \Pr(A|m_L))}{\Pr(A|m_H)}\theta_L, \tag{8}$$

where $\Pr(A|m_i) = \theta_i + x_i$, with x_i as the distortion in the decision rule. This expression contains the main insights regarding the agent's truth-telling constraint. First, the constraint can be binding only for the lower-quality alternative, as the gain comes from increasing the probability of acceptance. Second, even if the agent has a temptation to misrepresent only the low-quality alternative, the principal may optimally alter her decision rule for both alternatives.

Because the temptation to lie arises from the incremental increase in the acceptance probability, $\frac{\Pr(A|m_H) - \Pr(A|m_L)}{\Pr(A|m_H)}$, the first means through which the principal will manage the constraint is to increase the acceptance probability when the agent sends a weak recommendation. This increase in $\Pr(A|m_L)$ will both increase the agent's continuation value $u_A(x_L, x_H)$ and relax the reneging temptation. The second means is through altering the acceptance probability following a strong recommendation, $\Pr(A|m_H)$. Here, however, the effects go in opposite directions: increasing the acceptance probability following a strong recommendation both increases the continuation value (relaxing the constraint) and increases the immediate gain to deviation (tightening the constraint). As a result, as we will see below, the equilibrium distortion may be in either direction. This observation is also relevant for the case of non-stationary equilibrium considered below. If the principal rewards the agent by increasing his future influence, such a change may increase his reneging temptation in the future and make the relationship harder to sustain.

For the principal, the incentive-compatibility constraint arises from the fact that by deviating from the decision rule, she is able to save the distortion x_i. Thus, for her, the reneging temptation arises for both recommendations as long as $x_i \neq 0$. This constraint can be written as

$$\max|x_i| \leq \frac{\delta_P}{1-\delta_P}\left(v_P(x_H, x_L) - \underline{v}_P\right). \tag{9}$$

We can thus write the principal's maximization problem as

$$\min_{x_H, x_L} \left(p x_H^2 + (1-p) x_L^2\right)$$

$$\text{s.t.} \quad \begin{array}{l} \frac{\delta_A}{1-\delta_A}\left[u_A(x_L, x_H) - \underline{u}_A\right] \geq \frac{((\theta_H - \theta_L) + (x_H - x_L))}{(\theta_H + x_H)}\theta_L \\ \max|x_i| \leq \frac{\delta_P}{1-\delta_P}\left(v_P(x_H, x_L) - \underline{v}_P\right). \end{array}$$

Now, once the first-best can no longer be achieved, the agent's IC constraint will always be binding. If it is not binding, we could always decrease x_L, increasing the principal's payoff. So I will first consider the solution when the principal's IC is not binding (e.g. if $\delta_P \to 1$), and then consider how the solution changes once the principal's IC becomes a concern.

The logic behind the solution is easiest to illustrate graphically, as done in Figure 3. First, the principal's indifference curves are ellipses, and the solution will always have $x_L \geq 0$ to moderate the incentives to exaggerate. The relevant segments of the principal's indifference curves are denoted by the light grey dashed lines growing away from the origin, with the principal's preferred point at the origin. Second, examination of the agent's IC constraint shows that it will generally have two solutions for each x_L (indeed, the solution is given by the quadratic formula). The reason for this result is that altering the acceptance probability has the tradeoff mentioned above: increasing x_H increases the value of the relationship but also the gain from an immediate deviation. Thus, we can satisfy the IC constraint either by promising a high continuation value, achieved by positive x_H (we are biasing the high recommendation in the agent's favor), or by limiting the incentives to exaggerate by restricting the agent's influence when making a strong recommendation (negative x_H).

The agent's patience level then infuences the efficiency of these two avenues. Not surprisingly, when the agent is patient, promising a high future value is more efficient, while when the agent is impatient, providing an immediate reduction in the value of deviating is more efficient. This tradeoff is illustrated by the dark dashed lines in the Figure, where increasing the agent's impatience causes the IC constraints to "fan out." An important focal point is provided by the solution

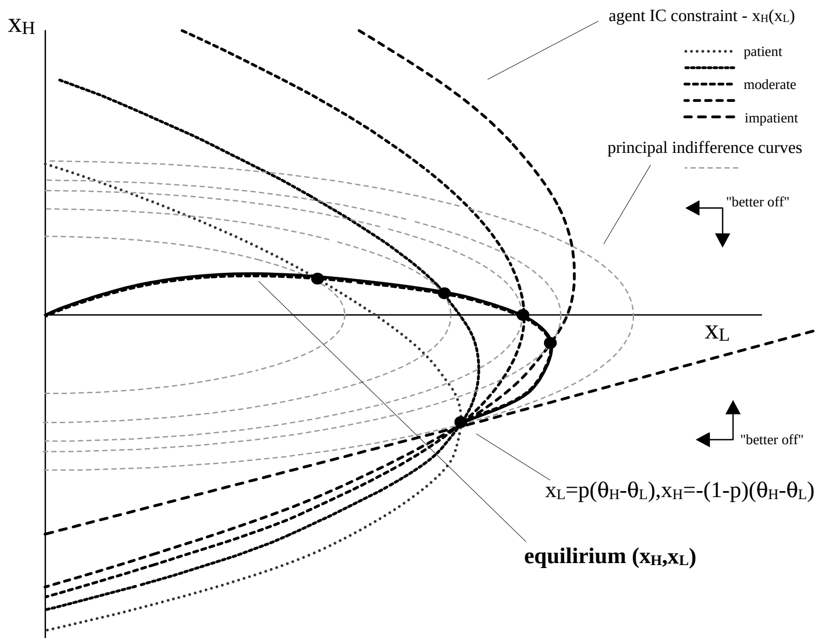

Figure 3: Deriving the equilibrium distortions

$(x_L = p(\theta_H - \theta_L), x_H = -(1-p)(\theta_H - \theta_L))$, which satisfies the agent's IC constraint for all patience levels. Indeed, at this point $\theta_H + x_H = \theta_L + x_L = E(\theta)$ and the principal ignores all information provided to her, so the agent's patience level is irrelevant.

The solution for any patience level, δ_A is then standard. Find the tangency point between the agent's IC constraint and principal's indifference curve, indicating the indifference curve that is closest to the origin. As we change the patience level, we trace out the $(x_L(\delta_A), x_H(\delta_A))$ function as illustrated. In particular, for high patience levels, the optimal solution involves favoring the agent for both moderate projects and high-quality projects. The agent is sufficiently patient that it is optimal to reward him with continued high influence. However, as the agent becomes increasingly impatient, we need to increase the influence of the agent for moderate project to sustain IC, but we start to decrease his influence for the high-quality proposals to limit the incentives to exaggerate. After a point, the cost of maintaining the IC constraint for the moderate projects becomes so high that the principal actually starts to discriminate against the high-quality projects and once such discrimination begins, she will eventually start also decreasing the favoritism for the medium-quality projects. As the agent becomes fully myopic, the only solution sustainable is the one that ignores the information presented.[3]

Next, note that this solution only satisfies the agent's IC constraint. As the agent becomes increasingly impatient, the distortions needed to maintain incentive-compatibility grow in size and

[3] As illustrated, the solution is a smooth curve. However, an element of the solution is that both the agent's IC constraint and the principal's indifference curves are concave to the origin, and we cannot rule out the possibility that the principal's indifference curve is not always smoothly contained inside the agent's truth-telling constraint in the vicinity of $x_H = 0$. In this case, there would be a discrete jump from a strictly positive x_H to a strictly negative x_H, with an associated jump in x_L. I have not, however, managed to find such an exception.

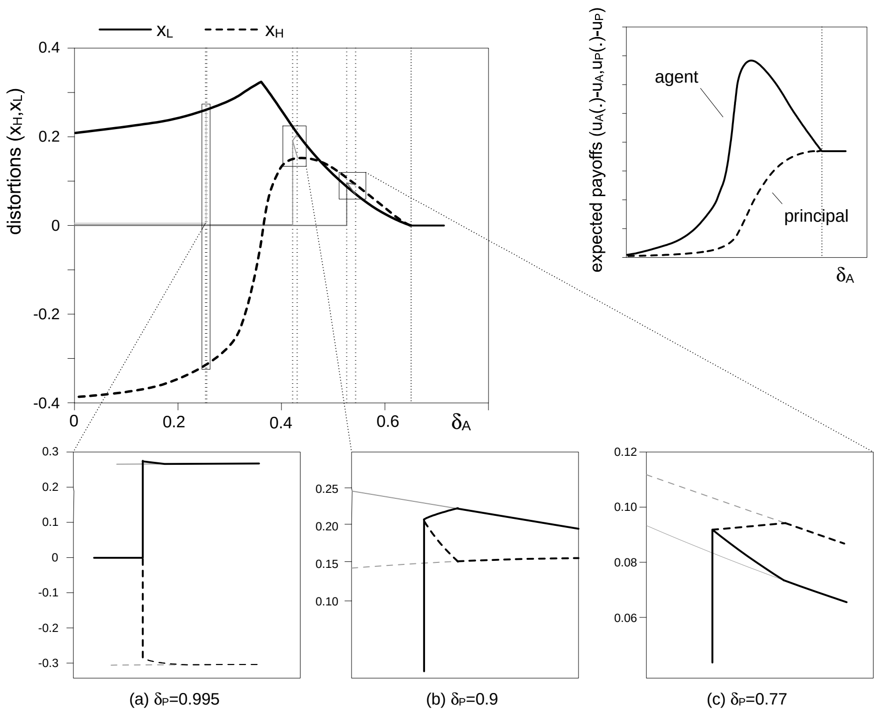

Figure 4: Optimal distortions and expected payoffs under stationary equilibrium - $\theta_L = 0.2, \theta_H = 0.8$ and $p = 0.35$.

the principal may become tempted to deviate herself. The principal's IC constraint then puts a cap on the distortion that can be introduced to the decision rule. The consequences of this cap are illustrated in Figure 4, which plots the unconstrained solution together with the constrained solution under different patience levels by the principal.

Three potential scenarios can arise, depending on the levels of (x_H, x_L) at the time the principal's constraint becomes binding. First, suppose that $x_H > x_L > 0$. Then, to satisfy the principal's constraint we must lower x_H and to continue to satisfy the agent's constraint we must increase x_L. However, as this is a less efficient means of satisfying the agent's constraint, the value of the relationship decreases, decreasing the maximal distortion sustainable. As the agent becomes increasingly impatient, the thresholds converge and the relationship becomes unsustainable – supporting further truth-telling by the agent would require increasing at least one of the thresholds, which is not possible without violating the principal's constraint.

The same logic applies in the other two cases, but the direction of adjustments changes. First, suppose that $x_L > x_H$, with $|x_H| < |x_L|$. In this case, the principal needs to decrease x_L to satisfy her own constraint, which increases the agent's incentives to exaggerate and requires an increase in the continuation value through x_H to compensate. Second, suppose that $x_L > 0 > x_H$ with $|x_H| > |x_L|$. In this case, the principal must increase x_H, which increases the incentives to exaggerate and must be compensated for by increasing x_L as well. Note that this scenario can arise only when $(1-p)(\theta_H - \theta_L)$ is sufficiently large so that the distortion can actually be sufficiently negative.

The basic results regarding the stationary equilibrium are summarized in the following proposition:

Proposition 2 ***Optimal stationary strategies:***
(i) The optimal acceptance rule of the principal is characterized by $x_L(\delta_A, \delta_P) \geq 0$ *and* $x_H(\delta_A, \delta_P) \gtreqless 0$ *with* $x_H(\delta_A, \delta_P) \leq 0$ *iff* $\delta_A \leq \underline{\delta}_A$.
(ii) The agent's relational influence (stage-game payoff) is non-monotone in δ_A. *It is increasing in* δ_A *for* $\delta_A \leq \widetilde{\delta}_A$ *and decreasing for* $\delta_A > \widetilde{\delta}_A$.
(iii) The bias in favor of medicore projects is monotone increasing in the impatience of the agent: $\frac{d(\Pr(A|m_H)/\Pr(A|m_L))}{d\delta_A} > 0$.

Proof. Part (i) follows directly from the analysis above. Part (ii) follows from the behavior of x_L and x_H. As the agent becomes impatient so that the first-best is no longer sustainable, he is initially rewarded with a higher continuation value by biasing the decision rules in his favor, but as he becomes increasingly impatient, the principal begins to discriminate against the better proposals, eroding the influence and eventually converging to the no-information benchmark. Part (iii) follows from the observation that the efficient distortion is proportional to the states (which would leave relative probability unchanged), and the optimal decision rule is increasingly distorted away from this to satisfy the truth-telling constraint of the agent. ■

The summarize, the basic insights from the analysis of the stationary equilibrium are as follows. First, to limit the incentives to exaggerate, the optimal decision rule always exhibits corporate socialism, in the sense that mediocre projects have a relatively higher likelihood of acceptance relative to the first-best outcome. This arises whether the high-quality projects are also favored or not. Indeed, the acceptance probability for the low quality may be higher than the first-best threshold for an average project. Second, the overall relative influence of the agent is non-monotone in his patience. This non-monotonicity followed from the result that increasing the likelihood of accepting high-quality projects both increased the continuation value to the agent (relaxing the truth-telling constraint) and increased the immediate gain from misleading the principipal (tightening the truth-telling constraint). When the agent is sufficiently patient, the first effect dominates and it is optimal to reward the agent with above-first-best relational influence to maintain truth-telling, whereas when the agent becomes more impatient, the temptation for the agent to abuse that high level of influence becomes too high and it becomes optimal to limit the agent's influence below the first-best level. In terms of project financing, this implies that the overall level of investment may be both above and below the first-best level.

6 Non-stationary strategies

While the distortions in the current influence of the agent can be used to sustain the relationship, a stationary policy fails to take into account the possibility of using changes in the future influence of the agent to achieve the same. The basic logic is relatively simple. First, to increase the attractiveness of truth-telling, the principal will reward the agent with additional future influence whenever he admits that his alternative is mediocre. For example, the dean of faculty may not hire a given job candidate or the CEO may choose not to finance a given project given lukewarm support by the department or division, but promises priority treatment the following year. This use of future influence brings two additional benefits. First, it allows the principal to smooth out the reward for truth-telling over multiple periods: instead of settling up immediately, the reward is spread over time by giving the agent higher expected level of influence in the future. Second, it brings about an important efficiency gain: since the agent values the implementation of high-quality projects more than low-quality projects, partially delaying the reward is valuable. Instead of having a low-quality project implemented today, he receives a promise of favorable treatment tomorrow, when he may have a high-quality project available.

In addition to rewarding the agent for the admission of mediocre projects, the principal will also optimally punish the agent in some settings. In particular, when the agent makes a strong recommendation and the proposal is still rejected, the truthfulness of that proposal is never learned and this limits the cost of deviation to the agent. To counter this, it will be optimal to follow such a rejected recommendation with a decrease in future influence. This result is akin to the punishment phases in games of imperfectly observed actions, such as the triggering of price wars in Green and Porter (1984) and lower effort levels by the agent in Li and Matouschek (2013).

To see these features in more detail, index the current influence of the agent by i. Then, following the equilibrium strategy (truth-telling) gives the agent an expected payoff of

$$\Pr(A|L,i)\,\theta_L + \delta_A \left[\Pr(A|L,i)\,V_A(A|L,i) + (1-\Pr(A|L,i))\,V_A(R|L,i)\right], \tag{10}$$

where L indicates the lower recommendation and A indicates an acceptance (and R rejection). Similarly, if the agent lies, then his continuation payoff is given by

$$\Pr(A|H,i)\,\theta_L + \delta_A \left[\Pr(A|H,i)\,V_A^{dev} + (1-\Pr(A|H,i))\,V_A(R|H,i)\right]. \tag{11}$$

Now, as the first preliminary observation, recall that the principal's loss in the stage-game is given by $px_{H,i}^2 + (1-p)\,x_{L,i}^2$ and is thus convex in the biases. Thus, to provide a given continuation value $[\Pr(A|L,i)\,V_A(A|L,i) + (1-\Pr(A|L,i))\,V_A(R|L,i)]$ at minimum cost to the principal, it is best to set $V_A(A|L,i) = V_A(R|L,i)$ as a way to minimize the cost of distortions. Then, the IC constraint simplifies to

$$\frac{\left(V_A(L,i) - V_A^{dev}\right) - \Pr(R|H,i)\left(V_A(R|H,i) - V_A^{dev}\right)}{(1-\Pr(R|H,i))} \geq \frac{1}{\delta_A}\left(1 - \frac{\Pr(A|L,i)}{\Pr(A|H,i)}\right)\theta_L. \tag{12}$$

In the case of the stationary equilibrium considered above, we had $V_A(L,i) = V_A(R|H,i)$, given by $\frac{1}{1-\delta_A} u_A(x_{H,i}, x_{L,i})$. Now, by setting $V_A(L,i) > \frac{1}{1-\delta_A} u_A(x_{H,i}, x_{L,i}) > V_A(R|H,i)$, we are able to relax the IC constraint for the agent and improve the current payoff by lowering the immediate thresholds. This improvement is achieved by switching to a more favorable equilibrium in the future following an admittance of a mediocre alternative, while punishing the agent if the principal chooses not to implement the alternative even when argued to be great. These are the two basic distortions discussed above.[4]

Now, going forward it will be more convenient to deal with the net payoffs, so define $U_j(.) = V_j(.) - V_j^{dev}$ as the normalized payoff, which allows us to write the truth-telling constraint as

$$U_A(L,i) - \Pr(R|H,i)\, U_A(R|H,i) \geq \frac{1}{\delta_A}\left(\Pr(A|H,i) - \Pr(A|L,i)\right)\theta_L. \tag{13}$$

The second constraint that the solution needs to satisfy is the promise-keeping constraint, which states that the expected payoff promised to the agent at the beginning of the period, $U_A(i)$, must equal the expected payoff from the stage game and the resulting continuation payoffs. In other words, we have

$$U_A(i) = (u_A(i) - \underline{u}_A) + \delta_A\left[(1-p)\, U_A(L,i) + p\left(\Pr(R|H,i)\, U_A(R|H,i) + \Pr(A|H,i)\, U_A(A|H,i)\right)\right]. \tag{14}$$

Now, since the continuation value following an accepted high recommendation plays no role in the truth-telling constraint, we minimize the variance in the continuation values by setting $U_A(A|H,i) = U_A(i)$, so that the agent's influence is unchanged following this outcome. Then, we can combine Equations 12 and 14 to yield the two key constraints that the changes in future influence must satisfy:

$$\Delta U_A(R|H,i) = \frac{U_A(i)\left(1 - \delta_A\left(1 - (1-p)\Pr(A|H,i)\right)\right) + (1-\delta_A)\, V_A^{dev} - E(\theta)\Pr(A|H,i)}{\delta_A \Pr(R|H,i)} \tag{15}$$

and

$$\Delta U_A(L,i) = \frac{U_A(i)\left[1 - \delta_A(1 + p\Pr(A|H,i))\right] + (1-\delta_A)\, V_A^{dev} - \Pr(A|H,i)E(\theta) + \left(\Pr(A|H,i) - \Pr(A|L,i)\right)\theta_L}{\delta_A}, \tag{16}$$

where $\Delta U_A(R|H,i) = U_A(R|H,i) - U_A(i) < 0$ and $\Delta U_A(L,i) = U_A(L,i) - U_A(i) > 0$ are the

[4] Note that if the constraint is binding in one state, it needs to be binding in all states. To see this, recall from Lemma 1 that the efficient means of delivering a given stage-game payoff $v_A(x_{H,i}, x_{L,i})$ to the agent involved a proportional distortion, $x_{j,i} = \alpha_i \theta_j$. But under a proportional distortion, RHS of Equation 12 is constant. Thus, if the IC constraint is not satisfied under the first-best solution of $\alpha_i = 0$ and $V(L,i) = V(R|H,i) = V(i)$, the principal must bear the cost of disorting either the LHS by introducing $V(L,i) > V(i) > V(R|H,i)$ or by distorting RHS by using an inefficient stage-game policy for delivering the same stage-game payoff. If the IC constraint is not binding, the principal could always do better by either reducing the distortion in the future by lowering the spread $V(L,i) - V(R|H,i)$ or by lowering the stage-game distortion, until the IC constraint is satisfied with equality.

changes in the continuation value. The principal's problem is then to maximize

$$U_P(i)(1-\delta_P) = (u_P(i) - \underline{u}_P) + \delta_P\left[(1-p)\Delta U_P(L,i) + p\Pr(R|H,i)\Delta U_P(R|H,i)\right] \tag{17}$$

subject to equations 15 and 16 at each stage of the game, together with the choice of the initial promise.

Now, the key element needed to solve the model completely would be to characterize the self-generating payoff set. Unfortunately, the richness of the principal's action space makes characterizing the set challenging.[5] We can, however, use the principle of optimality to characterize the basic tradeoff between settling up now through changes in the immediate acceptance probabilities and settling up later through changes in the continuation values without fully needing to characterize the frontier. Having considered the basic tradeoffs further, I then provide a numeric solution to a simplified three-state variant of the model.

6.1 Characterizing the dynamics

To consider the dynamics of the optimal relationship, we can make two preliminary observations. First, because there is only one-sided asymmetric information, the optimal contract can utilize payoffs on the Pareto frontier. Second, because the stage-game losses are convex in the magnitude of the distortions, the Pareto frontier will be concave.

Using these two observations together with equations 15, 16 and 17, we can derive some further insights to the transition dynamics. This process of transitions is illustrated in Figure 5. Begin with the initial state i_0. Given the promised value and the distortions x_{H,i_0} and x_{L,i_0}, equations 15 and 16 pin down the future promises $\Delta U_A(R|H,i_0)$ and $\Delta U_A(L,i_0)$ needed to satisfy the initial promise in an incentive-compatible fashion. This pins down transitions to $U_A(i_1)$ or $U_A(i_2)$. The transitions needed, however, depend on the stage-game distortions. In particular, increasing $x_{L,i}$, by increasing the stage-game payoff right now, lowers the need for rewarding the agent in the future, lowering $\Delta U_A(L,i)$. Similarly, increasing $x_{H,i}$ lowers the need to reward the agent in the future. However, because an increase in $x_{H,i}$ also increases the reneging temptation, $\Delta U_A(L,i)$ cannot decrease as rapidly as with $x_{L,i}$. As a result, an increase in $x_{H,i}$ also requires an increase in the punishment following rejection, so that $\Delta U_A(R|H,i)$ must decrease as well.

For interior states, the solution is then characterized by the following first-order conditions. For $x_{L,i}$ we have

$$\frac{\partial u_P(i)}{\partial x_{L,i}} + \delta_P\left[(1-p)\frac{\partial \Delta U_P(L,i)}{\partial \Delta U_A(L,i)}\frac{\partial \Delta U_A(L,i)}{\partial x_{L,i}}\right] = 0 \Leftrightarrow x_{L,i} = -\frac{\delta_P}{\delta_A}\theta_L\frac{\partial \Delta U_P(L,i)}{\partial \Delta U_A(L,i)}. \tag{18}$$

[5] The concavity of the frontier implies that the optimal solution does not have the bang-bang property that would allow us to characterize the Pareto frontier simply through its extreme points.

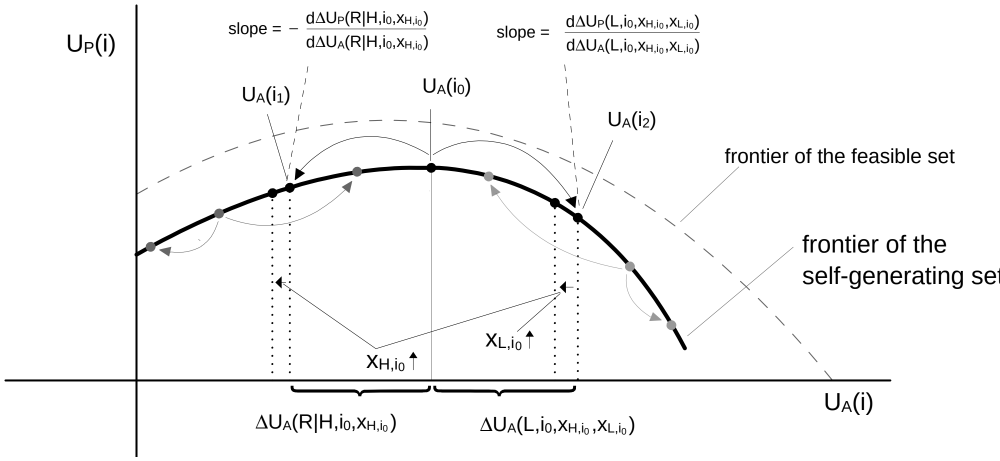

Figure 5: Illustrating the equilibrium dynamics

The basic tradeoff is then straightforward. Rewarding the agent immediately with an increase in acceptance probability costs the principal $x_{L,i}$ while rewarding the agent at the rate θ_L. Conversely, increasing the continuation value by $\Delta U_A(L,i)$ gives the agent $\delta_A \Delta U_A(L,i)$ while costing the principal $\delta_P \Delta U_P(L,i)$. Thus, we are simply equating the marginal cost/benefit ratios of the two channels.

The optimal distortion for the high-quality projects is slightly more complex, as $x_{H,i}$ influences both continuation values $\Delta U_A(L,i)$ and $\Delta U_A(R|H,i)$ and the likelihood of rejection, $\Pr(R|H,i)$. However, the first-order condition is still straightforward and given by

$$\begin{aligned}\frac{\partial u_P(i)}{\partial x_{H,i}} + \delta_P \left[(1-p)\frac{\partial \Delta U_P(L,i)}{\partial \Delta U_A(L,i)}\frac{\partial \Delta U_A(L,i)}{\partial x_{H,i}}\right] \\ -\delta_P p \Delta U_P(R|H,i) + \delta_P \left[p \Pr(R|H,i)\frac{\partial \Delta U_P(R|H,i)}{\partial \Delta U_A(R|H,i)}\frac{\partial \Delta U_A(R|H,i)}{\partial x_{H,i}}\right] = 0.\end{aligned} \tag{19}$$

The basic forces are thus as follows. First, increasing $x_{H,i}$ bears the basic cost (or benefit, if negative) of $-px_{H,i}$. Second, by rewarding the agent now, it allows the principal to lower the reward going forward, through $\frac{\partial \Delta U_A(L,i)}{\partial x_{H,i}} < 0$. Third, it lowers the likelihood that a good project is rejected, and the associated change in the continuation value, $\Delta U_P(R|H,i)$. Finally, to satisfy the promise-keeping constraint, it increases the punishment that needs to be imposed on the agent upon rejection. Using the first-order condition for $x_{L,i}$ and equations 15 and 16, we can write the above expression as

$$\begin{aligned}x_{H,i} = (1-p)\frac{x_{L,i}(\delta_A U_A(i) + (\theta_H - \theta_L))}{\theta_L} \\ -\delta_P \Delta U_P(R|H,i) - \frac{\delta_P}{\delta_A}\left[\frac{E(\theta) - (1-p\delta_A)\Delta U_A(i) - (1-\delta_A)V_A^{dev}}{\Pr(R|H,i)}\right]\frac{\partial \Delta U_P(R|H,i)}{\partial \Delta U_A(R|H,i)}.\end{aligned} \tag{20}$$

In other words, because of the link between $x_{L,i}$ and $\frac{\partial \Delta U_P(L,i)}{\partial \Delta U_A(L,i)}$, we can observe that the benefit of increasing $x_{H,i}$ translates to the ability to decrease $x_{L,i}$.

An important feature for both conditions is that both the slopes $\frac{\partial \Delta U_P(.)}{\partial \Delta U_A(.)}$ and the change in the principal's expected payoff following a rejection, $\Delta U_P(R|H,i)$, depend on the initial state, i. For states close to i_0, $\frac{\partial \Delta U_P(L,i)}{\partial \Delta U_A(L,i)}$ and $\frac{\partial \Delta U_P(R|H,i)}{\partial \Delta U_A(R|H,i)}$ are both negative - increasing the reward or punishment to the agent requires a lower payoff to the principal. Thus, equation 18 gives us $x_{L,i} > 0$ - some settling up occurs immediately to avoid needing to reward the agent too much in the future. Similarly, since the principal transitions to a worse state following a rejection of a high-quality alternative, $\Delta U_P(R|H,i) < 0$ and so all three components of equation 20 are positive, giving us $x_{H,i} > 0$.

But as we move further away from the status quo, things will be different. First, suppose the agent has built up a lot of relational influence. Then, the promise-keeping constraint requires that even a rejection keeps us above the status quo. In this case, increasing $x_{H,i}$ has a dual benefit - not only does it allow us to lower the need for additional (and increasingly expensive) future promises but it also allows us to *increase* the penalty from rejection, which in this case benefits the principal by bringing her closer to the status quo. Conversely, for very low states, increasing $x_{H,i}$ has a dual cost – not only does it increase the need for further penalty, which is increasingly costly as we move away from the status quo, but it also lowers the ability to return towards the status quo. Thus, the dynamics have an inherent force towards the status quo. The higher the current state, the higher the cost of additional future promises, leading to higher settling up now through $x_{L,i}$ and $x_{H,i}$, combined with high loss of influence following rejection, $\Delta U_A(R|H,i)$, and increasingly small addition of further influence, $\Delta U_A(L,i)$. Conversely, the lower the current state, the lower $x_{L,i}$ and $x_{H,i}$ to complete the punishment immediately and thus lower the need for further future punishment and create a faster return towards status quo.

Caveats: It should be noted that the above discussion is mainly descriptive. As mentioned above, to fully characterize the equilibrium would require us to solve for the Pareto frontier, which is computationally infeasible. Also, the characterization applies only when the principal is able to transition optimally both upward and downward. But both the rewards and penalties are bounded and the corner solutions don't allow for the same level of flexibility of determining the parameters optimally. Finally, not knowing the frontier prevents us from deriving further, more detailed comparative statics.

6.2 A three-state example

Given the difficulties in characterizing the optimal relational contract further, I will now illustrate some of the basic ideas using a three-state model. A solution with a highly patient principal is illustrated in Figure 6, where panels (i) and (ii) illustrate the distortions and transition probabilities. We can see that the basic logic of the stationary equilibrium continues to hold. The principal favors mediocre projects while favoring high-quality projects for intermediate discount rates while discriminating against them for low discount rates. Further, not surprisingly, the distortions are ranked, with the higher state having a higher acceptance probability for both types of projects.

More importantly, we can see that the distortions in the non-stationary setting tend to be lower than in the stationary setting, even for the extremal states. The reason is the ability to use the

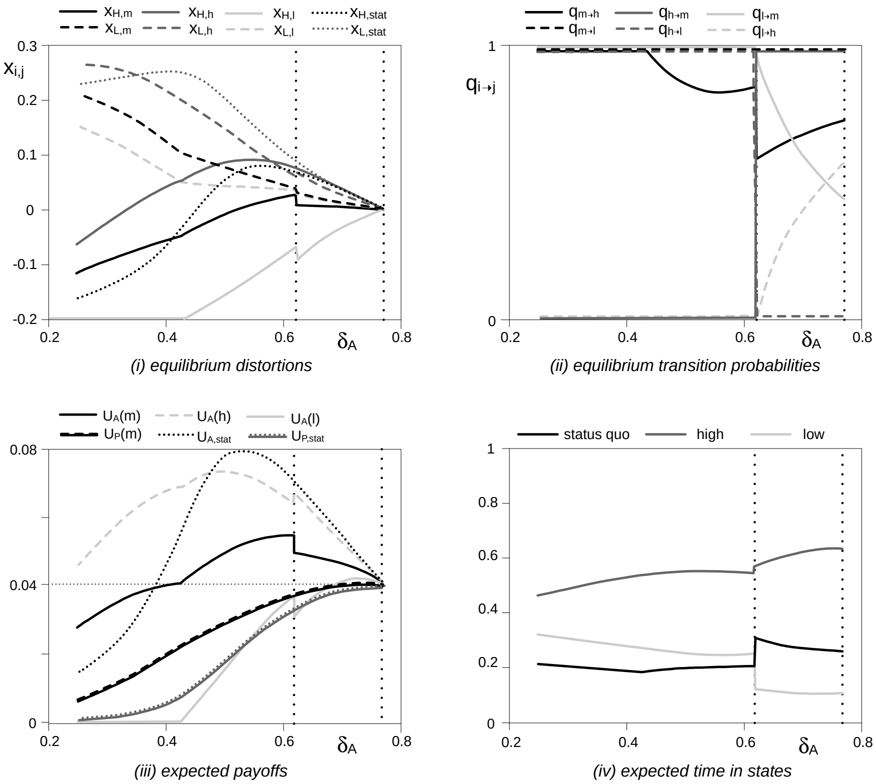

Figure 6: An example of the three-state solution, $\theta_L = 0.2, \theta_H = 0.6$ and $p = 0.5, \delta_P = 0.95$.

changes in the future influence instead of immediate settling up to manage the relationship. The benefit of this is most clearly illustrated in panel (iii), which plots the principal's expected payoff, together with the agent's continuation payoffs in the different states. For intermediate states, the changes in the continuation value allow us to limit the influence that we need to grant the agent – the agent's expected payoff even in the high state is below the agent's payoff in the stationary equilibrium. Correspondingly, for low levels of patience, the changes in the future influence substitute for need to limit the agent's influence, leading to a higher payoff for both parties.

The second element of interest is the transition probabilities themselves. Panel (iv) illustrates the net effect of the transition probabilities the cleanest by considering the expected time spent in the different states. Much like with the stage-game distortions, when the agent is moderately impatient, he values the relationship a lot and, as a result, the game spends a high expected amount of time in the high state. As the agent becomes increasingly impatient, the reneging temptation grows, and it becomes more efficient to penalize the agent, as the reneging temptation in the low state is lower. Thus, the the parties spend less time in the high state and an increasing amount of time in the low state.

In addition to these two basic themes, the solution exhibits two kinks. The first arises because of the three-state nature of the model. As a part of the increasing need to motivate the agent through penalties, increasing impatience of the agent lowers the transition probability up from the

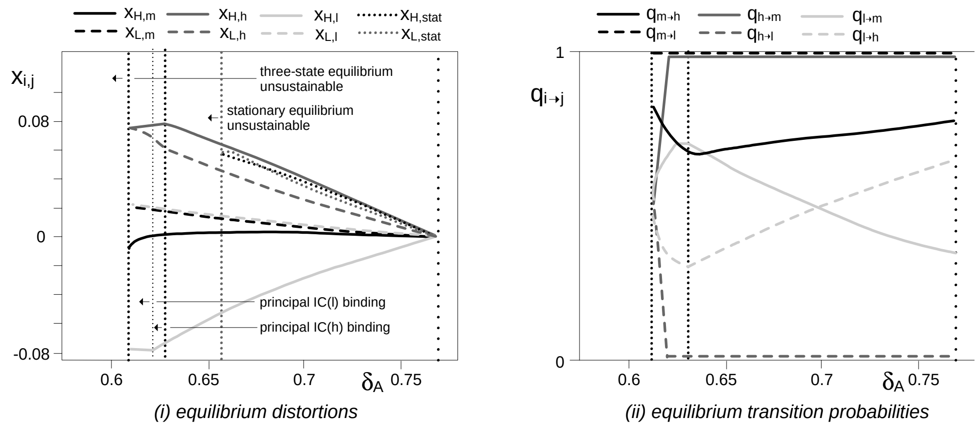

Figure 7: An example of the three-state solution, $\theta_L = 0.2, \theta_H = 0.6$ and $p = 0.5, \delta_P = 0.77$.

low state. To the right of the discontinuity, $q_{l \to h} > 0$ and the principal benefits from the lower transition probability by increasing the likelihood that the parties return to the status quo state instead of transitioning directly to the high state. At $q_{l \to h} = 0$ this benefit disappears and the value function exhibits a kink, with further reductions associated with loss to the principal. Thus, the probability settles on $q_{l \to m} = 1$ and the rest of the parameters adjust to take into account this change. Finally, when the agent is sufficiently impatient, it becomes optimal to push the agent all the way to his outside option in the low state - and since no further punishment of the agent is possible, the parameters again exhibit a kink at this point.

This solution illustrated the outcome under a highly patient principal. If the principal is impatient, then again the constraint on the incentive-compatibility of the disortion can become a binding constraint, and the mechanics are very similar to the stationary scenario, with the principal first adjusting the distortions to sustain incentive-compatibility, until the relationship becomes unsustainable. The only differences are that (i) because of the ability to sustain a higher value to the relationship with smaller distortions (panels (i) and (iii)), the relationship is sustainable longer, and (ii) transition probabilities provide an additional lever to satisfy both constraints simultaneously. One such solution is illustrated in Figure 7, which plots the solution under the non-stationary equilibrium, together with the stationary equilibrium, for a less patient principal. To highlight the role of the transition probabilities, note that as both high and low state constraints become binding, the optimal solution correspondingly increases the transition probability to that state to limit the distortion away from the promised continuation value.

7 Conclusion

I have illustrated how a decision-maker can use the repeated nature of a relationship and the manipulation of her acceptance rule to sustain the relationship with an agent. In a stationary equilibrium, to sustain truth-telling, the decision-maker implements a form of corporate socialism, where the decisions are biased in favor of mediocre alternatives to reward honesty. In a non-stationary equilibrium, the principal can further reward the agent through the allocation of future influence, where the admission of mediocre alternatives is rewarded with increased future influence while the rejections of high-quality alternatives are associated with the erosion of influence.

These results are, however, only a first pass towards a richer understanding of the evolution of influence in organizations. In terms of dynamics, work remains to be done on what is the truly optimal equilibrium. The analysis also assumed that the outcomes are perfectly observed. In reality, such observation is likely to be noisy, which leads to at least three complications. First, when the observations are noisy, the optimal equilibrium is likely to involve comparing the blocks of the empirical distribution of the recommendations to the predicted distribution under honesty. Second, if the arrival of information is probabilistic, then the realm of deviations becomes larger as an agent who deviates once will take into account the gains he can get from a continued deviation until he is caught. Third, the detection of the deviation may be dependent on the magnitude of the lie, which creates the question of optimal deviation and the possibility of needing to tolerate minor misrepresentations along the equilibrium path.

The analysis also ignored the possibility of transfers. In that regard, deep pockets would convert the problem into a stationary one as transfers could be used as signals of private information. If the transfers can flow only from the principal to the agent, then the basic qualitative features of the solution would be maintained, but large enough distortions would be supplemented by monetary transfers to limit the cost to the principal. Relatedly, use of performance pay creates the possibility of partial alignment of interests. In the extreme, making the agent the residual claimant would eliminate all conflicts and lead to the first-best outcome.

In terms of the timing and observability of information, I assumed that the cost of implementation is public and arrives only after the agent's recommendation. In general, hiding the information from the agent until his recommendation is optimal, but a fully private signal would convert the problem into one of two-sided private information and again necessitate history-dependent punishments by the agent if his advice is not followed often enough. Finally, an interesting question that remains open is how to manage a relationship with multiple agents who provide competing proposals or otherwise multiple relevant pieces of information in terms of the allocation of current and future influence among them.

References

[1] Abreu, Dilip, David Pearce and Ennio Stacchetti (1990), "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, 58(5), pp.1041-1063

[2] Athey, Susan and Kyle Bagwell (2001), "optimal collusion with private information," rand, 32(3), 428-65

[3] Athey, Susan, Kyle Bagwell and Chris Sanchirico (2004), "Collusion and Price Rigidity," Review of Economic Studies, 71, 317-349

[4] Alonso, Ricardo and Niko Matouschek (2007), "Relational Delegation," *The RAND Journal of Economics*, 38(4), 1070-1089

[5] Andrews, Isaiah and Daniel Barron (2014), "The Allocation of Future Business: Dynamic Relational Contracts with Multiple Agents," working paper, Harvard Society of Fellows and Northwestern Kellogg

[6] Baker, George, Robert Gibbons and Kevin Murphy (1999), "Informal Authority in Organizations," Journal of Law, Economics, and Organization, 15(1), pp. 56-73

[7] Barron, Daniel (2013), "Putting the Relationship First: Relational Contracts and Market Structure," working paper, Northwestern Kellogg

[8] Barron, Daniel (2014), "Attaining Efficiency with Imperfect Public Monitoring and One-Sided Markov Adverse Selection," working paper, Northwestern Kellogg

[9] Barron, Daniel and Michael Powell (2015), "Policies in Relational Contracts," working paper, Northwestern Kellogg

[10] Board, Simon (2011), "Relational Contracts and the Value of Loyalty," American Economic Review, 101, 3349-3367

[11] Calzolari, Giacomo and Giancarlo Spagnolo (2010), "Relational Contracts and Competitive Screening," CEPR Discussion Paper No. 7434

[12] Campbell, Arthur (2015), "Political Capital," working paper, Yale University

[13] Chakraborty, Archishman and Bilge Yilmaz (2013), "Authority, Consensus and Governance," working paper, Yeshiva University and University of Pennsylvania

[14] Crawford, Vincent and Joel Sobel (1982) "Strategic Information Transmission," *Econometrica, 50,* 1431-1451

[15] Garfagnini, Umberto, Marco Ottaviani and Peter Sorensen (2014), "Accept or Reject? An Organizational Perspective," *International Journal of Industrial Organization*, 34, pp. 66-74

[16] Guo, Yingni and Johannes Horner (2015), "Dynamic Mechanisms without Money," Cowles foundation discussion paper No. 1985

[17] Kolotilin, Anton and Hongyi Li (2015), "Relational Communication with Transfers," working paper, University of New South Wales

[18] Levin, Jonathan (2002), "Multilateral Contracting and the Employment Relationship," The Quarterly Journal of Economics, pp. 1075-1103

[19] Levin, Jonathan (2003), "Relational Incentive Contracts," American Economic Review, 93(3), 835-857

[20] Li, Jin and Niko Matouschek (2013), "Managing Conflicts in Relational Contracts," AER 103(6), pp.2328-2351

[21] Li, Jin, Niko Matouschek and Michael Powell (2015), "Power Dynamics in Organizations," working paper, Northwestern University

[22] Li, Zhouzheng, Heikki Rantakari and Huanxing Yang (2016) "Competitive Cheap Talk," Games and Economic Behavior, 96, pp. 65-89

[23] Lipnowski, Elliot and Joao Ramos (2015), "Repeated Delegation," working paper, New York University

[24] Malcomson, James (2015), "Relational Incentive Contracts with Private Information," University of Oxford Department of Economics Discussion Paper Series No. 634

[25] Padro i Miquel, Gerard and Pierre Yared (2012), "The Political Economy of Indirect Control," QJE 127, 947-1015

[26] Rantakari, Heikki (2015), "Managerial Influence and Organizational Performance," working paper, University of Rochester

[27] Rantakari, Heikki (2016), "Soliciting Advice: Active versus Passive Principals," forthcoming, Journal of Law, Economics, and Organization

[28] Renault, Jérôme, Eilon Solan and Nicolas Vieille (2013), "Dynamic sender-receiver games," Journal of Economic Theory, 148(2), pp. 502-534

A Proofs and derivations

A.1 Solving for the stationary equilibrium

To examine the properties of the solution (x_H, x_L), we can first write the IC constraint of the agent, which we know to be always binding, as

$$\left[\phi(\theta) + p\theta_H x_H + (1-p)\,\theta_L x_L\right] - y_A \tfrac{((\theta_H-\theta_L)+(x_H-x_L))}{(\theta_H+x_H)}\theta_L = 0,$$

where $\phi(\theta)$ is the value created by the information and $y_A = \frac{1-\delta_A}{\delta_A}$. The principal's goal is to minimize the loss caused by the distortion, $px_H^2 + (1-p)\,x_L^2$ subject to the IC constraint. Taking the Lagrangian of the problem, we can write the two first-order conditions for the solution as

$$\begin{aligned} &x_H : px_H - \lambda\left(p\theta_H - y_A \tfrac{(\theta_L+x_L)}{(\theta_H+x_H)^2}\theta_L\right) = 0 \\ &x_L : (1-p)\,x_L - \lambda\left((1-p)\,\theta_L + y_A \tfrac{1}{(\theta_H+x_H)}\theta_L\right) = 0, \end{aligned}$$

which gives us that the optimal (x_H, x_L) must satisfy

$$\frac{px_H}{\left(p\theta_H - y_A \frac{(\theta_L+x_L)}{(\theta_H+x_H)^2}\theta_L\right)} = \frac{(1-p)x_L}{\left((1-p)\theta_L + y_A \frac{1}{(\theta_H+x_H)}\theta_L\right)}.$$

This expression allows us to learn some of the basic features of the solution. First, RHS is always positive and thus we have that $x_L \geq 0$. Second, the left-hand side may be positive or negative, depending on the magnitude of y_A. When the agent is sufficiently patient, y_A is small and thus $x_H > 0$. Conversely, if the agent is sufficiently impatient, $x_H < 0$. Further, we can use the agent's IC constraint to solve explicitly for the solution for x_H. The IC constraint yields a quadratic equation in x_H, the solution to which is given by

$$x_H\left(x_L, \phi(\theta)\right) = \frac{-\left[\phi(\theta)+(1-p)\theta_L x_L + p\theta_H^2 - y_A\theta_L\right] \pm \sqrt{\left[\phi(\theta)+(1-p)\theta_L x_L + p\theta_H^2 - y_A\theta_L\right]^2 - 4p\theta_H\left[\left[\phi(\theta)+(1-p)\theta_L x_L\right]\theta_H - y_A(\theta_H-\theta_L-x_L)\theta_L\right]}}{2p\theta_H}.$$

While deriving the rest of the solution analytically is infeasible, we can obtain three points of the solution. First, first-best is obtained when the solution satisfies $x_H = x_L = 0$, which gives the threshold level of patience from the analysis as $y_A = \frac{\phi(\theta)\theta_H}{(\theta_H-\theta_L)\theta_L}$. Second as the limit when the agent becomes infinitely impatient, the IC constraint converges to $x_L = ((\theta_H - \theta_L) + x_H)$ and the solution is $(x_H = -(1-p)(\theta_H - \theta_L), x_L = p(\theta_H - \theta_L))$, so that $\theta_H + x_H = E(\theta) = \theta_L + x_L$. Finally, at the point of $x_H = 0$, we know the LHS of the expression above must be ill-defined, giving $\frac{p\theta_H^3}{y_A\theta_L} - \theta_L = x_L$.

To formally consider the monotonicity of the solution, we can make two observations. First, for a given x_L, the principal's indifference curves become flatter at given $x_H(x_L)$ as we reach lower indifference curves (further from the origin). To see this, note that the indifference curves are given by

$$px_H^2 + (1-p)\,x_L^2 = L_P.$$

Thus, the slope at any given point is given by

$$\frac{dx_H}{dx_L} = -\frac{(1-p)x_L}{px_H}.$$

But for x_H to lie on the indifference curve, we have that $x_H = \pm\sqrt{\frac{L_P-(1-p)x_L^2}{p}}$, so that the slope is

$$\frac{dx_H}{dx_L} = -\frac{(1-p)x_L}{px_H} = \pm\frac{(1-p)x_L}{p\sqrt{\frac{L_P-(1-p)x_L^2}{p}}},$$

which is flattening in the loss, L_P. Further, $\frac{d^2x_H}{dx_L^2} = -\frac{(1-p)}{px_H}$, so that the function is concave to the origin. In other words, $\left|\frac{dx_H}{dx_L}\right|$ is increasing in x_L. Second, for the agent's incentive-compatibility constraint, we can establish that for each indifference curve, $\left|\frac{dx_H}{dx_L}\right|$ is also increasing in x_L.

Next, considering the effect of impatience, we can establish from the agent's IC constraint that $\frac{dx_H}{d\delta_A} > 0$ for $\frac{dx_H}{dx_L} < 0$ and $\frac{dx_H}{d\delta_A} < 0$ for $\frac{dx_H}{dx_L} > 0$ over the relevant range. This is the "fanning out" of the IC constraint discussed in the text. Further, we have that $\frac{d^2x_H(x_L)}{dx_L d\delta_A} \leq 0$, so that the IC constraint becomes pointwise steeper in the region of $\frac{dx_H}{dx_L} > 0$ while becoming less positive for $\frac{dx_H}{dx_L} < 0$. Then, the solution is immediate. Consider first solution in the region of $\frac{dx_H}{dx_L} > 0$. Then, an increase in the impatience causes an outward-shift of the IC constraint, pushing the principal to a worse indifference curve. This curve is pointwise flatter than the original curve. Correspondingly, the IC constraint is pointwise steeper. Thus, the new tangency must lie to the right, with x_L higher than the original solution. Similarly, consider the region of $\frac{dx_H}{dx_L} > 0$. Now, both the principal and the agent have pointwise flatter indifference curves, and the direction of the change in x_L is indeterminate, but x_H must decrease. To see this, begin with a fully impatient agent, with a tangency point at the uninformative solution. A decrease in impatience increases the slope $\frac{dx_H}{dx_L}$ for the IC constraint of the agent, rotating it around the point of indifference, allowing the principal to obtain a lower indifference curve, which means that the new tangency point must lie to the northeast of the original tangency point. As the agent becomes more patient, the tangency point must continue to travel north, but will revert to moving west once the agent is sufficiently patient.

To establish the first part of the properties of the agent's IC constraint, note that we can use the implicit function theorem to derive

$$\frac{dx_H}{dy_A} = \frac{\theta_L(\theta_H+x_H)(((\theta_H-\theta_L)+(x_H-x_L)))}{\left(p\theta_H(\theta_H+x_H)^2-y_A(\theta_L+x_L)\theta_L\right)} \qquad \textit{and} \qquad \frac{dx_H}{dx_L} = -\frac{\theta_L\theta_H+x_H(y_A+(1-p)(\theta_H+x_H))}{\left(p\theta_H(\theta_H+x_H)^2-y_A(\theta_L+x_L)\theta_L\right)},$$

which establishes the opposite signs of $\frac{dx_H}{dy_A}$ and $\frac{dx_H}{dx_L}$ and thus the "fanning" out of the feasible solutions. Deriving $\frac{d^2x_H(x_L)}{dx_L d\delta_A} \leq 0$ is more cumbersome, since it requires us taking into account the change in the location of the IC constraint. As a first step, it follows immediately from the IC constraint that $\frac{d^2x_H}{dx_L d\delta_A} < 0$. However, this does not take into account the fact that a change in δ_A also changes x_H. To this end, we can use the exact solution for $x_H\left(x_L\right)$ from above. The derivative is analytically tedious, but the relevant part simplifies to

$$\frac{d^2x_H(x_L)}{dx_L d\delta_A} \leq 0 \Leftrightarrow$$

$$\left[p\theta_H\left(\theta_H-\left(\frac{(\theta_H+x_H)}{((\theta_H+x_H)-(\theta_L+x_L))}\right)x_H\right)-\left(1+\frac{(\theta_H+x_H)}{((\theta_H-\theta_L)+(x_H-x_L))}\right)\phi(\theta)\right]$$
$$\leq(1-p)\,\theta_L x_L\left[1+\left(\frac{(\theta_H+x_H)}{((\theta_H-\theta_L)+(x_H-x_L))}\right)\right],$$

where $\phi(\theta)$ is the value of information. The expression is immediately satisfied whenever p is small enough, as the condition can be violated only when $p\theta_H^2$ is sufficiently high. Unfortunately we cannot verify the whole parameter space analytically (the condition can be violated for some generic (x_L, x_H) so we need to take into account the constraint that $x_H(x_L, \delta_A)$ must hold, which we can verify numerically).

A.2 Equations for the non-stationary equilibrium and the three-state model

The truth-telling constraint of the agent is given by

$$\frac{V_A(L,i)-\Pr(R|H,i)V_A(R|H,i)}{(1-\Pr(R|H,i))}\geq\frac{1}{\delta_A}\left(1-\frac{\Pr(A|L,i)}{\Pr(A|H,i)}\right)\theta_L+V_A^{dev},$$

while the promise-keeping constraint is given by

$$V_A(i)=u_A(i)+\delta_A\left[(1-p)\,V_A(L,i)+p\left(\Pr(R|H,i)\,V_A(R|H,i)+\Pr(A|H,i)\,V_A(i)\right)\right],$$

which we can rewrite as

$$V_A(i)\left(1-\delta_i\left(1+p\Pr(A|H,i)\right)\right)=u_A(i)-\delta_A p\left[V_A(L,i)-\Pr(R|H,i)\,V_A(R|H,i)\right].$$

From the IC constraint we get that

$$\left[V_A(L,i)-\Pr(R|H,i)\,V_A(R|H,i)\right]=\frac{1}{\delta_A}\left(\Pr(A|H,i)-\Pr(A|L,i)\right)\theta_L+\Pr(A|H,i)\,V_A^{dev},$$

which we can substitute in the promise-keeping constraint to yield

$$V_A(i)=\frac{u_A(i)+\delta_A V_A(L,i)-p(\Pr(A|H,i)-\Pr(A|L,i))\theta_L-\delta_A p\Pr(A|H,i)V_A^{dev}}{[1-\delta_A p\Pr(A|H,i)]}.$$

Next, we can decompose the stage-game payoff as

$$u_A(i)=\phi(\theta)+p\theta_H x_H(i)+(1-p)\,\theta_L x_L(i)+(1-\delta_A)\,V_A^{dev},$$

where $(1-\delta_A)\,V^{dev}$ is the payoff from an uninformed choice while $\phi(\theta)=p(1-p)(\theta_H-\theta_L)^2$ is the expected value of the agent's information. We can substitute this into the expression, which then simplifies the expression to

$$\left(V_A(i)-V_A^{dev}\right)=\frac{\phi(\theta)+p(x_{H,i}-\theta_L)(\theta_H-\theta_L)+\theta_L x_{L,i}+\delta_A\left(V_A(L,i)-V_A^{dev}\right)}{[1-\delta_A p\Pr(A|H,i)]}.$$

As an additional simplification to the expression, we can add and substract $p\theta_H(\theta_H - \theta_L)$ and θ_L^2 from the expression, which gives us

$$\left(V_A(i) - V_A^{dev}\right) = \frac{\phi(\theta) + p\Pr(A|H,i)(\theta_H - \theta_L) - p\theta_H^2 - (1-p)\theta_L^2 + \Pr(A|L,i)\theta_L + \delta_A\left(V_A(L,i) - V_A^{dev}\right)}{[1 - \delta_A p\Pr(A|H,i)]},$$

and finally noting that $p\theta_H^2 + (1-p)\theta_L^2$ is the value realized by the agent in the first-best scenario and can be decomposed as $\phi(\theta) + (1-\delta_A)V_A^{dev}$, we finally get

$$\left(V_A(i) - V_A^{dev}\right) = \frac{(E(\theta) - \theta_L)\Pr(A|H,i) + \Pr(A|L,i)\theta_L + \delta_A V_A(L,i) - V_A^{dev}}{[1 - \delta_A p\Pr(A|H,i)]}.$$

Then, we can rearrange the expression, using $U_A(i) = V_A(i) - V_A^{dev}$, to

$$\Delta U_A(L,i) = \frac{U_A(i)[1 - \delta_A(1 + p\Pr(A|H,i))] + (1-\delta_A)V_A^{dev} - E(\theta)\Pr(A|H,i) + (\Pr(A|H,i) - \Pr(A|L,i))\theta_L}{\delta_A}$$

Alternatively, we can take the promise-keeping constraint and add and substract $(1-p)V_A(R|H,i)$, which allows us to obtain a different factorization of the same expression as

$$\left(V_A(i) - V_A^{dev}\right) = \frac{\phi(\theta) + x_{H,i}(p\theta_H + (1-p)\theta_L) + (1-p)(\theta_H - \theta_L)\theta_L + \delta_A\Pr(R|H,i)\left[V_A(R|H,i) - V_A^{dev}\right]}{(1 - \delta_A p\Pr(A|H,i))},$$

and using a manipulation similar to the above, we can write this as

$$\left(V_A(i) - V_A^{dev}\right) = \frac{E(\theta)\Pr(A|H,i) - \delta_A\Pr(A|H,i)\left(V_A(R|H,i) - V_A^{dev}\right) + \delta_A V_A(R|H,i) - V_A^{dev}}{(1 - \delta_A p\Pr(A|H,i))},$$

which we can rearrage to

$$\Delta U_A(R|H,i) = \frac{U_A(i)(1 - \delta_A(1 - (1-p)\Pr(A|H,i))) + (1-\delta_A)V_A^{dev} - E(\theta)\Pr(A|H,i)}{\delta_A\Pr(R|H,i)}.$$

Given these two expressions, we can then consider the principal maximizing her payoff,

$$U_P(i)(1-\delta_P) = (u_P(i) - \underline{u}_P) + \delta_P\left(p\Pr(R|H,i)\Delta U_P(R|H,i) + (1-p)\Delta U_P(L,i)\right),$$

subject to the above two constraints. The first-order conditions are then immediate. First, for the choice of $x_{L,i}$ gives

$$\frac{dU_P(i)}{dx_{L,i}} = \frac{\partial u_P(i)}{\partial x_{L,i}} + \delta_P(1-p)\frac{\partial \Delta U_P(L,i)}{\partial \Delta U_A(L,i)}\frac{d\Delta U_A(L,i)}{dx_{L,i}} = 0,$$

and from above we get $\frac{d\Delta U_A(L,i)}{dx_{L,i}} = -\frac{\theta_L}{\delta_A}$, which simplifies the expression to

$$x_{L,i} = -\frac{\delta_P}{\delta_A}\frac{\partial \Delta U_P(L,i)}{\partial \Delta U_A(L,i)}\theta_L.$$

Similarly, for the distortions for the high project, we get

$$\frac{dU_P(i)}{dx_{H,i}} = \frac{\partial u_P(i)}{\partial x_{H,i}} + \delta_P(1-p)\frac{\partial \Delta U_P(L,i)}{\partial \Delta U_A(L,i)}\frac{d\Delta U_A(L,i)}{dx_{H,i}}$$

$$-\delta_P\left(p\Delta U_P(R|H,i)\right)+\delta_P p\Pr(R|H,i)\frac{\partial\Delta U_P(R|H,i)}{\partial\Delta U_A(R|H,i)}\frac{d\Delta U_A(R|H,i)}{dx_{H,i}}=0.$$

From above, we get

$$\begin{array}{l} x_{L,i}=-\frac{\delta_P}{\delta_A}\frac{\partial\Delta U_P(L,i)}{\partial\Delta U_A(L,i)}\theta_L\Leftrightarrow-\frac{x_{L,i}\delta_A}{\delta_P\theta_L}=\frac{\partial\Delta U_P(L,i)}{\partial\Delta U_A(L,i)},\\ \frac{d\Delta U_A(L,i)}{dx_{H,i}}=-\frac{\delta_A pU_A(i)+(E(\theta)-\theta_L)}{\delta_A}<0\\ \frac{d\Delta U_A(R|H,i)}{dx_{L,i}}=\frac{U_A(i)(1-p\delta_A)+(1-\delta_A)V^{dev}-E(\theta)}{\delta_A\Pr(R|H,i)^2}<0, \end{array}$$

where the last result follows from the following logic. For the transition to be valid, it must be that

$$\Delta U_A\left(R|H,i\right)=\frac{U_A(i)(1-\delta_A(1-(1-p)\Pr(A|H,i)))+(1-\delta_A)V^{dev}-E(\theta)\Pr(A|H,i)}{\delta_A\Pr(R|H,i)}<0,$$

which holds only as long as

$$\left[U_A\left(i\right)\delta_A(1-p)-E\left(\theta\right)\right]<-\frac{V_A(i)(1-\delta_A)}{\Pr(A|H,i)},$$

so that

$$\max\left(\frac{d\Delta U_A(R|H,i)}{dx_{H,i}}\right)=\frac{(1-\delta_A)V_A(i)-\frac{V_A(i)(1-\delta_A)}{\Pr(A|H,i)}}{\delta_A\Pr(R|H,i)^2}=-\frac{(1-\delta_A)V_A(i)}{\delta_A\Pr(R|H,i)\Pr(A|H,i)}<0.$$

Substituting in the components, we get

$$\begin{array}{l} x_{H,i}=(1-p)\frac{x_{L,i}(\delta_A U_A(i)+(\theta_H-\theta_L))}{\theta_L}-\delta_P\Delta U_P(R|H,i)\\ \quad-\delta_P\frac{\partial\Delta U_P(R|H,i)}{\partial\Delta U_A(R|H,i)}\frac{E(\theta)-U_A(i)(1-p\delta_A)-(1-\delta_A)V^{dev}}{\delta_A\Pr(R|H,i)^2}. \end{array}$$

A.2.1 Three-state model

Principal's expected payoffs: To solve for the principal's expected payoff that we are trying to maximize, we can use the value functions and rewrite them to yield

$$\begin{array}{l} U_P\left(h\right)=\frac{u_h+\delta_P p\widetilde{q}_{h\to m}U_P(m)+\delta_P p\widetilde{q}_{h\to l}U_P(l)}{[(1-\delta_P)+\delta_P p(\widetilde{q}_{h\to m}+\widetilde{q}_{h\to l})]}\\ U_P\left(l\right)=\frac{u_l+(1-p)\delta_P[q_{l\to m}U_P(m)+q_{l\to h}U_P(h)]}{[(1-\delta_P)+(1-p)\delta_P(q_{l\to m}+q_{l\to h})]}\\ U_P\left(m\right)=\frac{u_m+\delta_P q_{m\to h}(1-p)U_P(h)+\delta_P p\widetilde{q}_{m\to l}U_P(l)}{[(1-\delta_P)+\delta_P q_{m\to h}(1-p)+\delta_P p\widetilde{q}_{m\to l}]}, \end{array}$$

where $\widetilde{q}_{j\to i}=\Pr(R|H,i)q_{j\to i}$ is the expected transition probability. Next, define constants

$$\begin{array}{l} \alpha_h=\left[(1-\delta_P)+\delta_P p\left(\widetilde{q}_{h\to m}+\widetilde{q}_{h\to l}\right)\right]\\ \alpha_l=\left[(1-\delta_P)+(1-p)\,\delta_P\left(q_{l\to m}+q_{l\to h}\right)\right]\\ \alpha_m=\left[(1-\delta_P)+\delta_P q_{m\to h}\left(1-p\right)+\delta_P p\widetilde{q}_{m\to l}\right] \end{array}$$

as the denominators of the three expressions, we can then solve the three equations simultaneously to yield the value functions as a function of the transition probabilities and the stage-game payoffs as

$$U_P\left(l\right)=\frac{u_l(\alpha_m\alpha_h-\beta_{hm})+(1-p)\delta_P(q_{l\to m}\alpha_h+\delta_P p q_{l\to h}\widetilde{q}_{h\to m})u_m+(1-p)\delta_P(q_{l\to h}\alpha_m+(1-p)\delta_P q_{l\to m}q_{m\to h})u_h}{\alpha_l\alpha_m\alpha_h-(1-p)\delta_P^2 p(\delta_P p\widetilde{q}_{h\to m}\widetilde{q}_{m\to l}q_{l\to h}+\delta_P(1-p)\widetilde{q}_{h\to l}q_{l\to m}q_{m\to h}+\alpha_h q_{l\to m}\widetilde{q}_{m\to l}+\alpha_l\widetilde{q}_{h\to m}q_{m\to h}+\alpha_m\widetilde{q}_{h\to l}q_{l\to h})}$$

$$U_P\left(m\right)=\frac{u_m[\alpha_l\alpha_h-\beta_{hl}]+(1-p)\delta_P(q_{m\to h}\alpha_l+\delta_P p\widetilde{q}_{m\to l}q_{l\to h})u_h+\delta_P p(\widetilde{q}_{m\to l}\alpha_h+(1-p)\delta_P q_{m\to h}\widetilde{q}_{h\to l})u_l}{\alpha_l\alpha_m\alpha_h-(1-p)\delta_P^2 p(\delta_P p\widetilde{q}_{h\to m}\widetilde{q}_{m\to l}q_{l\to h}+\delta_P(1-p)\widetilde{q}_{h\to l}q_{l\to m}q_{m\to h}+\alpha_h q_{l\to m}\widetilde{q}_{m\to l}+\alpha_l\widetilde{q}_{h\to m}q_{m\to h}+\alpha_m\widetilde{q}_{h\to l}q_{l\to h})}$$

$$U_P\left(h\right)=\frac{u_h[\alpha_l\alpha_m-\beta_{ml}]+\delta_P p(\widetilde{q}_{h\to l}\alpha_m+\delta_P p\widetilde{q}_{h\to m}\widetilde{q}_{m\to l})u_l+\delta_P p(\widetilde{q}_{h\to m}\alpha_l+(1-p)\delta_P\widetilde{q}_{h\to l}q_{l\to m})u_m}{\alpha_l\alpha_m\alpha_h-(1-p)\delta_P^2 p(\delta_P p\widetilde{q}_{h\to m}\widetilde{q}_{m\to l}q_{l\to h}+\delta_P(1-p)\widetilde{q}_{h\to l}q_{l\to m}q_{m\to h}+\alpha_h q_{l\to m}\widetilde{q}_{m\to l}+\alpha_l\widetilde{q}_{h\to m}q_{m\to h}+\alpha_m\widetilde{q}_{h\to l}q_{l\to h})},$$

where

$$\begin{aligned}
\beta_{hl} &= \delta_P^2 p\left(1-p\right)\widetilde{q}_{h\to l}q_{l\to h}\\
\beta_{hm} &= \delta_P^2 p\left(1-p\right)\widetilde{q}_{h\to m}q_{m\to h}\\
\beta_{ml} &= \delta_P^2 p\left(1-p\right)\widetilde{q}_{m\to l}q_{l\to m}.
\end{aligned}$$

State-specific constraints: In similar fashion, we can use the constraints on the continuation payoffs to explicitly consider the three states.

High state: In the high state, no further transition upwards is possible, so $V_A\left(L,h\right)=V_A\left(h\right).$ For the punishment, the continuation value is given by

$$V_A\left(R|H,i\right)=V_A(h)+q_{h\to m}\left(V_A(m)-V_A(h)\right)+q_{h\to l}\left(V_A(l)-V_A(h)\right). \tag{21}$$

This allows us to write the constraints as

$$\Pr\left(A|L,h\right)=\frac{\left(V_A\left(h\right)-V_A^{dev}\right)\left[1-\delta_A\left(1+p\Pr\left(A|H,h\right)\right)\right]+\left(1-\delta_A\right)V_A^{dev}-\left(E\left(\theta\right)-\theta_L\right)\Pr\left(A|H,h\right)}{\theta_L}, \tag{22}$$

and

$$\Delta V_A\left(R|H,h\right)=\frac{\left(V_A\left(h\right)-V_A^{dev}\right)\left(1-\delta_A\left(1-\left(1-p\right)\Pr\left(A|H,h\right)\right)\right)+\left(1-\delta_A\right)V_A^{dev}-E\left(\theta\right)\Pr\left(A|H,h\right)}{\delta_A\Pr\left(R|H,h\right)}, \tag{23}$$

Note that because further transition upwards is not possible, the RHS of both equations is a function of only $x_{H,h}$. Thus, for the largest state, the optimization problem reduces into a single-variable problem. Finally, we can write the principal's payoff as

$$V_P\left(h\right)\left(1-\delta_P\right)=u_P\left(h\right)+\delta_P p\Pr(R|H,h)\Delta V_P(R|H,h). \tag{24}$$

Finally, noting that because of the convexity of the losses to the principal when moving away from the best solution, the transition probabilities satisfy either $q_{h\to m} < 1, q_{h\to l} = 0$ or $q_{h\to l} > 0, q_{h\to m} = 1 - q_{h\to l}$, we get

$$\Delta V_P(R|H,h) = \begin{cases} q_{h\to m}\left(V_P(m) - V_P(h)\right) & \text{for } q_{h\to l} = 0 \\ \left(V_P(m) - V_P(h)\right) + q_{h\to l}\left(V_P(l) - V_P(m)\right) & \text{for } q_{h\to l} > 0 \end{cases}. \tag{25}$$

Then using equation 23 we get the solution for the transition probabilities as

$$\begin{array}{ll} q_{h\to m} = \frac{\left(V_A(h) - V_A^{dev}\right)(1-\delta_A(1-(1-p)\Pr(A|H,h))) + (1-\delta_A)V_A^{dev} - E(\theta)\Pr(A|H,h)}{\delta_A \Pr(R|H,h)(V_A(m) - V_A(h))} & \text{for } q_{h\to l} = 0 \\ q_{h\to l} = \frac{\left(V_A(h) - V_A^{dev}\right)(1-\delta_A(1-(1-p)\Pr(A|H,h))) + (1-\delta_A)V_A^{dev} - E(\theta)\Pr(A|H,h) - (V_A(m) - V_A(h))}{\delta_A \Pr(R|H,h)(V_A(l) - V_A(m))} & \text{for } q_{h\to l} > 0 \end{array}. \tag{26}$$

Given the initial promises $V_A(m), V_A(h), V_A(l)$, these equations define a range $\left[\underline{x}_{H,h}, \overline{x}_{H,h}\right]$ and the associated $x_{L,h}, q_{h\to m}$ and $q_{h\to l}$ for which a solution exists and we can maximize over.

Low state: For the low state, we can follow a similar logic. Now, no further transitions downwards are possible and, as a result, $V_A(R|H,l) = V_A(l)$ while the continuation value following a weak recommendation is given by

$$V_A(L|l) = V_A(l) + q_{l\to m}\left(V_A(m) - V_A(l)\right) + q_{l\to h}\left(V_A(h) - V_A(l)\right), \tag{27}$$

and the promise-keeping and truth-telling constraints simplify to

$$\Pr(A|H,l) = \frac{V_A(l)(1-\delta_i)}{E(\theta) - \left(V_A(l) - V_A^{dev}\right)\delta_A(1-p)} \tag{28}$$

and

$$\Delta V_A(L,l) = \frac{\left(\Pr(A|H,l) - \Pr(A|L,l)\right)\theta_L - \left(V_A(l) - V_A^{dev}\right)\delta_A \Pr(A|H,l)}{\delta_A}. \tag{29}$$

Since no further movement downwards is possible, there is a unique distortion $x_{H,l}$ that satisfies both constraints. Second, given the deterministic $\Pr(A|H,l)$, there is now a one-to-one relationship between $V_A(L,l)$ and $\Pr(A|L,l)$. As with the high state, the transition probabilities then become

$$\begin{array}{ll} q_{l\to m} = \frac{(\Pr(A|H,i) - \Pr(A|L,i))\theta_L - \delta_A \Pr(A|H,i)\left(V_A(l) - V_A^{dev}\right)}{\delta_A (V_A(m) - V_A(l))} & \text{for } q_{l\to h} = 0 \\ q_{h\to l} = \frac{(\Pr(A|H,i) - \Pr(A|L,i))\theta_L - \delta_A \Pr(A|H,i)\left(V_A(l) - V_A^{dev}\right) - \delta_A (V(m) - V(l))}{\delta_A (V_A(h) - V_A(m))} & \text{for } q_{l\to h} > 0 \end{array}. \tag{30}$$

Finally, the principal's value function needs to satisfy.

$$V_P(l)(1-\delta_P) = u_P(l) + (1-p)\delta_P \Delta V_P(L|l) \tag{31}$$

And the constraints define now a feasible range $\left[\underline{x}_{L,l}, \overline{x}_{L,l}\right]$ and the associated $x_{H,l}, q_{l\to m}$ and $q_{l\to h}$ for which a solution exists for the initial promises, and that we can maximize over.

The only qualitative difference to the continuous-state case under both the high and low states is that the value functions exhibit a kink at $q_{i\to m} = 1$, where for lower total transition probabilities, the principal benefits from an increased transition probability (as it approaches the status quo), whereas for higher transition probabilities, the principal is hurt by transitioning away from the status quo. Thus, the solution will exhibit potential discontinuities when the solution is around this point.

Status quo state: Finally, we can consider the principal's maximization problem for the status quo state. Take $V_A(m)$ to be the utility promised by the principal in the first stage of the game. Now, transitioning both up and down is possible, and we can write the principal's expected payoff as

$$V_P(m)(1-\delta_P) = u_P(m) + \delta_P\left(p\Pr(R|H,m)\Delta V_P(R|H,m) + (1-p)\Delta V_P(L|m)\right), \tag{32}$$

where $\Delta V_P(R|H,m) = q_{m\to l}(V_P(l) - V_P(m))$ and $\Delta V_P(L|m) = q_{m\to h}(V_P(h) - V_P(m))$. The relevant constraints are then given by

$$q_{m\to l} = \frac{\left(V_A(m) - V_A^{dev}\right)(1-\delta_A(1-(1-p)\Pr(A|H,m))) + (1-\delta_A)V_A^{dev} - E(\theta)\Pr(A|H,m)}{\delta_A \Pr(R|H,m)(V_A(l) - V_A(m))}. \tag{33}$$

Given the relationship between $\Delta V_A(R|H,m)$ and $\Pr(A|H,m)$, now there is a similar relationship between $\Delta V_A(L|m)$ and $\Pr(A|L,m)$. In particular, we have

$$q_{m\to h} = \frac{(1-\delta_A)V_A(m) - \delta_A p\Pr(A|H,i)\left(V_A(m) - V_A^{dev}\right) - (E(\theta) - \theta_L)\Pr(A|H,m) - \Pr(A|L,m)\theta_L}{\delta_A(V_A(h) - V_A(m))}. \tag{34}$$

The principal's problem in this stage then boils down to maximizing 32 subject to constraints 33 and 34, which can be viewed as a two-variable maximization problem with respect to $x_{H,m}$ and $x_{L,m}$.

Numerical optimization: The problem itself is solved numerically with Matlab. Using the principle of optimality, the solver consists of two stages. First, I optimize the optimal distortions in the three states following the initial promises of continuation values through a best-response iteration. Second, knowing how to optimally deliver the promised continuation values, I optimize over the initial promises $V(l), V(m)$ and $V(h)$.

B Continuous-state model

Consider the same setup as in the analysis, but assume now that the agent's return is drawn from a uniform distribution on [0,1]. Now, the acceptance rule becomes a function $\Pr(A|c, m_i)$. My focus is on the stage-game distortions, but to allow for us to consider the role of future influence as well, I allow the principal to promise a transition to a state with higher influence for the agent following the admittance of a low-quality alternative. I will, however, leave the state itself exogenous. Finally, I will focus only on the truth-telling constraint of the agent.

B.1 Attaining the first-best

Let $\Pr(A|m_i) = \int_c \Pr(A|m_i, c)\, dF_P(c)\, dc$ denote the expected probability of acceptance following a given message. Then, assuming that we can attain the first-best, the continuation equilibrium is stationary, and telling the truth gives an expected payoff of

$$\Pr(A|\theta_i)\,\theta_i + \delta_A V_A^{cont}. \tag{35}$$

Similarly, if the agent chooses to lie, his payoff is given by

$$\Pr\left(A|\widetilde{\theta}_i\right)\theta_i + \delta_A\left[\left(1 - \Pr\left(A|\widetilde{\theta}_i\right)\right)V_A^{cont} + \Pr\left(A|\widetilde{\theta}_i\right)V_A^{dev}\right]. \tag{36}$$

The optimal decision rule involves the principal accepting the project whenever $\theta_i \geq c$, so that $\Pr\left(A|\widetilde{\theta}_i\right) = F_P\left(\widetilde{\theta}_i\right) = \widetilde{\theta}_i$. Then, the truth-telling constraint reduces to, for all θ_i and $\widetilde{\theta}_i$, to

$$\left(V_A^{cont} - V_A^{dev}\right) \geq \frac{1}{\delta_A}\left(1 - \frac{F_P(\theta_i)}{F_P\left(\widetilde{\theta}_i\right)}\right)\theta_i. \tag{37}$$

Two observations follow. First, the optimal deviation for the agent is to always send the maximal message, which, given the assumed shared support and the first-best decision rule, leads to acceptance with probability 1. Intuitively, if it is optimal for the agent to risk burning his reputation, he should go all in. Second, the deviation temptation is maximized for interior θ_i. Since $F_P\left(\widetilde{\theta}_i\right) = 1$, the deviation temptation is maximized to θ_i for which

$$\theta_i = \frac{1 - F_P(\theta_i)}{f_P(\theta_i)}. \tag{38}$$

In the case of the uniform distribution, this yields $\theta_i = \frac{1}{2}$. Intuitively, when the state is low enough, the return to pushing for its acceptance is not worth the destruction of the reputation when caught lying. Similarly, when the state is high enough, then the project is relatively likely to be accepted even without exaggeration, and the small improvement in the probability of acceptance is not worth risking the reputation. It is this non-monotonicity of the deviation temptation that is the main difference to the two-state variant.

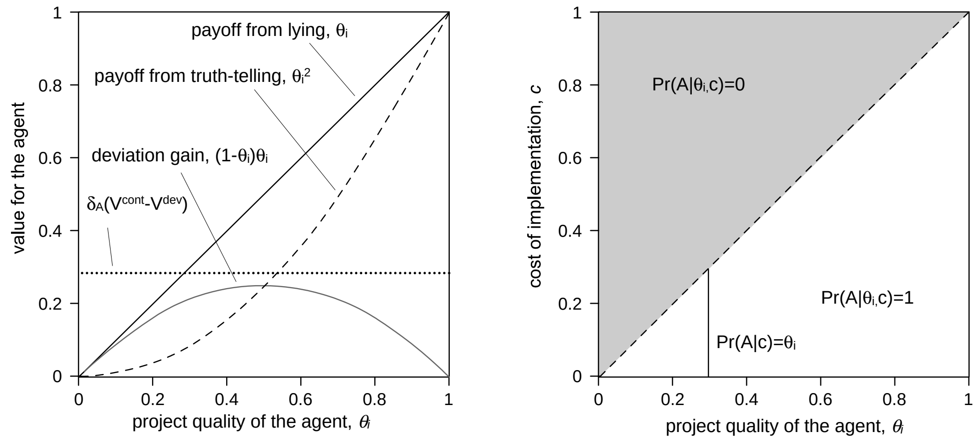

(i) Agent's payoff structure under the first-best decision rule

(ii) First-best decision rule

Figure 8: Sustainable truth-telling under the first-best decision rule.

This first-best solution is illustrated in Figure 8 for the uniform distribution. Panel (ii) plots the first-best decision rule. Simply, the agent tells the truth and the principal implements the agent's proposal if the cost is below the value of the project. Conditional on the decision rule, panel (i) illustrates the state-contingent payoffs of the agent, highlighting how the deviation temptation is maximized at $\theta_i = \frac{1}{2}$. Thus, if the truth-telling constraint is satisfied for $\theta_i = \frac{1}{2}$, it is slack for all other states. A quick valuation of the constraint reveals that truth-telling and first-best decision rule are incentive-compatible as long as $\delta_A \geq \frac{3}{4}$.

B.2 Managing the truth-telling constraint

If $\delta_A < \frac{3}{4}$, then the first-best decision rule is unable to sustain truth-telling by the agent. To analyze how we can optimally modify the constraint, note that we may condition the continuation payoff of the agent on both what his report is and whether the report is accepted or rejected. Thus, the expected payoff from truth-telling is given by

$$\Pr\left(A|\theta_i\right)\theta_i + \delta_A\left[\Pr\left(A|\theta_i\right)V_A^{cont}\left(\theta_i, A\right) + \left(1 - \Pr\left(A|\theta_i\right)V_A^{cont}\left(\theta_i, R\right)\right)\right], \tag{39}$$

where $V_A^{cont}\left(\theta_i, j\right)$ is the (potentially message-dependent) continuation value of the agent if he tells the truth. Conversely, if the agent chooses to lie, his payoff is given by

$$\Pr\left(A|\widetilde{\theta}_i\right)\theta_i + \delta_A\left[\left(1 - \Pr\left(A|\widetilde{\theta}_i\right)\right)V_A^{cont}\left(\widetilde{\theta}_i, R\right) + \Pr\left(A|\widetilde{\theta}_i\right)V_A^{dev}\right], \tag{40}$$

where $\widetilde{\theta}_i$ is the announcement of the agent, leading to a continuation payoff $V_A^{cont}\left(\widetilde{\theta}_i, R\right)$ if the principal chooses against the agent's proposal and thus the lie goes undetected, while resulting in payoff V_A^{dev} if the proposal is accepted and the agent is thus caught in a lie.

To capture the two-state structure of the payoffs, I index the states by $\underline{\alpha}$ and $\overline{\alpha}$, with the transition probability determined by $q_j(\theta_i)$,where $j \in \{A(\text{ccept}), R(\text{eject})\}$. Thus, we can write the continuation value in each case as

$$V_A^{cont}(\theta_i, j) = q_j(\theta_i)\,\Delta V_A + V_A(\underline{\alpha}), \tag{41}$$

where $\Delta V_A = V_A(\overline{\alpha}) - V_A(\underline{\alpha})$, the gain in the continuation value for the agent if the principal transitions to the more favorable decision rule. Then, taking equations 39 and 40, the decision rule and the continuation payoffs are incentive-compatible, if

$$\begin{aligned}&\Pr(A|\theta_i)\left[V_A^{cont}(\theta_i, A) - V_A^{cont}(\theta_i, R)\right] + \left[V_A^{cont}(\theta_i, R) - V_A^{cont}\left(\widetilde{\theta}_i, R\right)\right]\\ &+\Pr\left(A|\widetilde{\theta}_i\right)\left(V_A^{cont}\left(\widetilde{\theta}_i, R\right) - V_A^{dev}\right) \geq \tfrac{1}{\delta_A}\left(\Pr\left(A|\widetilde{\theta}_i\right) - \Pr(A|\theta_i)\right)\theta_i.\end{aligned}$$

First, it is immediate from the expression that for states high enough, the constraint is satisfied for any continuous acceptance rule.[6] Thus, following a deviation, the continuation value will have $q_R\left(\widetilde{\theta}_i\right) = 0$. Thus, the above simplifies to

$$\Pr(A|\theta_i)\left(q_A(\theta_i) - q_R(\theta_i)\right)\Delta V_A + q_R(\theta_i)\,\Delta V_A + \overline{p}\left(V_A(\underline{\alpha}) - V_A^{dev}\right) \geq \frac{1}{\delta_A}\left(\overline{p} - \Pr(A|\theta_i)\right)\theta_i. \tag{42}$$

Now, we face the question of how to optimally differentiate between $q_A(\theta_i)$ and $q_R(\theta_i)$. Note that because the agent cares only about the expected probability, as long as the optimal deviation is to make the strongest recommendation in favor of adoption, we can set $q_A(\theta_i) = q_R(\theta_i)$ for this stage of the analysis.

The caveat to this is that since the agent's continuation value is no longer independent of his report and, indeed, will be maximized for intermediate reports (to encourage truth-telling when it is most tempting to lie), if the change in continuation value is sufficiently high, while the intermediate types need to be discouraged from pretending to be high types, low types may now have an incentive to report being an intermediate type, in hopes of having their choice rejected and instead rewarded with future influence.

To counter this deviation, the principal could load more of the reward to actually accepting the proposal, so that $q_A(\theta_i) > q_R(\theta_i)$. This will, however, increase the maximal reward that the principal must give the agent in a given state (since the probabilities are capped from above by one, we need to increase ΔV_A, which is costly to the principal). Further, the bigger the distortion, the less attractive it becomes for the principal to honor the promise. Since the goal is not to provide a complete analysis of the framework, I assume that this constraint is not binding and let $q_A(\theta_i) = q_R(\theta_i)$.

[6] Continuous decision rules will turn out to be optimal when dealing with the agent's IC constraint. This changes once we introduce the principal's IC constraint as well.

Within these parameters, equation 42 helps us to identify four different means of managing the constraint, which I will consider in a sequence below.

1. General favoritism: The first immediate means of managing the agent's incentives to deviate is to increase the agent's influence in the stage game. In other words, we can increase $E\left[\Pr\left(A|\theta_i\right)\right]$ to increase the agent's expected payoff and thus the value of the relationship. The question is then what is the most efficient means of increasing the agent's stage-game payoff, ignoring any other constraints, and the solution is given by the following lemma:

Lemma 3 *The least-cost means of increasing the agent's continuation value is to use a linear distortion in the acceptance rule, with the principal implementing the project whenever* $c \leq \overline{c}\left(\theta_i\right)$, *where*

$$\overline{c}\left(\theta_i\right) = \left(1+\alpha\right)\theta_i, \alpha \geq 0. \tag{43}$$

Proof. See Appendix B.4.1 ∎

In other words, because the agent's payoff is proportional to the value of the idea, the least distortionary means of delivering a given utility to the agent is a proportional distortion, where the principal is willing to implement projects at a loss, and where the proportional loss is capped by α. As α increases, the agent is better off (and the principal is worse off).

2. Discrimination against the best proposals: The second observation that follows from equation 42 is that the agent's optimal deviation is going to be to the proposal that maximizes the probability of acceptance.[7] Further, we can see that the deviation temptation is increasing in the maximum probability of acceptance. Thus, the second means of satisfying the agent's truth-telling constraint is to lower the maximum probability of acceptance, which means that the agent's best proposals will be discriminated against. Indeed, the least-cost way of doing this is to introduce a cap $\overline{\overline{\theta}}_i$, where all proposals in $\left[\overline{\overline{\theta}}_i, 1\right]$ are accepted with a common probability $\Pr\left(A|\overline{\overline{\theta}}_i\right) = \overline{p} < 1$. Further, given part (1), the threshold and the probability of acceptance are implicitly related through $\overline{\overline{\theta}}_i\left(1+\alpha\right) = \overline{p}\left(\alpha\right)$.

To understand why the maximal acceptance probability plays a role even if the deviation temptation is maximized for intermediate types is as follows. When the agent chooses whether to lie or not, he doesn't yet know whether the lie is actually needed to induce acceptance. The lie is needed only when the cost of implementation is high enough, while truth would be enough when the cost is low enough. When choosing whether to mislead the principal, the agent is balancing these two forces. Now, if we lower the maximal probability of acceptance, so that the principal never implements the agent's proposal when her cost is high enough, the relative efficiency of the lie is lowered: it is more likely that telling the truth would have been sufficient to induce acceptance. Thus, exaggeration becomes less attractive.

[7] This result is immediate when the continuation payoff is not state-dependent. When the continuation payoff is state-dependent, we need to verify this to be the case.

In relation to the two-state framework of the main analysis, the level of distortion for the better project in the two-state framework blends these two effects (increasing the continuation value and decreasing the reneging temptation), where the direction of forces works in the opposite directions. In the continuous-state model, they become more decoupled as the principal can discrimate against the best projects while still providing favoritism towards above-average quality projects.

3. Leniency towards average proposals: As noted above, general favoritism is expensive, in the sense that it increases the agent's payoff for all types, even for those that would be willing to be truthful under a less favorable decision rule. An alternative means is then to increase the acceptance probability only for those types for whom it is directly needed. Following the logic from above, the deviation temptation is always maximized for intermediate states. Then, the concavity of the deviation tempation implies that if the truth-telling constraint of equation 42 is violated at some $\widehat{\theta}_i$, then there exist bounds $\underline{\theta} < \widehat{\theta}_i < \overline{\theta}$ such that the truth-telling constraint is just satisfied, with the constraint being slack outside the range and binding on the interior. As shown in Appendix B.4.2, we can then write, for the interior states, the probability of acceptance needed in equilibrium to induce truth-telling as a function of $\underline{\theta}$ as

$$\Pr\left(A|\theta_i\right) = \overline{p}\left(\alpha\right) - \left(\overline{p}\left(\alpha\right) - \left(1+\alpha\right)\underline{\theta}\right)\frac{\underline{\theta}}{\theta_i} \geq \left(1+\alpha\right)\theta_i, \tag{44}$$

where the thresholds solve

$$\left\{\underline{\theta}, \overline{\theta}\right\} = \frac{\overline{p}\left(\alpha\right)}{2\left(1+\alpha\right)}\left(1 \pm X\right), \qquad X = \sqrt{1 - \frac{4\left(1+\alpha\right)\delta_A \Delta V^{cont}}{\overline{p}\left(\alpha\right)^2}}, \tag{45}$$

where ΔV^{cont} is the expected change in continuation value for the agent if he decides to not be truthful.

4. Promise of future influence: Finally, instead of simply settling up today with the agent, the principal can promise the agent additional influence in the future. We capture this by the transition probability $q\left(\theta_i\right) \geq 0$ of switching to the more favorable state, $\overline{\alpha}$. For now, ignoring the whole design problem of $\Pr\left(A|\theta_i, \overline{\alpha}\right)$ and assuming that we have designed the state optimally, leading to the value $V_A\left(\overline{\alpha}\right)$, we can consider the tradeoff between settling now through $\Pr\left(A|\theta_i, \underline{\alpha}\right) > \left(1+\alpha\right)\theta_i$ or settling in the future through $q\left(\theta_i\right) > 0$. This tradeoff is resolved in the following Lemma:

Lemma 4 *Let* $\frac{\delta_P \Delta V_P}{\delta_A \Delta V_A}$ *to denote the cost-benefit ratio of switching from* $\underline{\alpha}$ *to* $\overline{\alpha}$ *equilibrium and* $\Pr\left(A|\theta_i\right)$ *the distortion needed to settle up using the current state alone. Then, the optimal way of settling up is*
(i) If $\Pr\left(A|\theta_i\right) < \left(1 + \frac{\delta_P \Delta V_P}{\delta_A \Delta V_A}\right)\theta_i$*, use the current decision only*
(ii) If $\Pr\left(A|\theta_i\right) > \left(1 + \frac{\delta_P \Delta V_P}{\delta_A \Delta V_A}\right)\theta_i$*, distort the current decision up to* $\overline{\Pr}\left(A|\theta_i\right) = \left(1 + \frac{\delta_P \Delta V_P}{\delta_A \Delta V_A}\right)\theta_i$*, and use* $q\left(\theta_i\right) > 0$ *to settle the rest.*

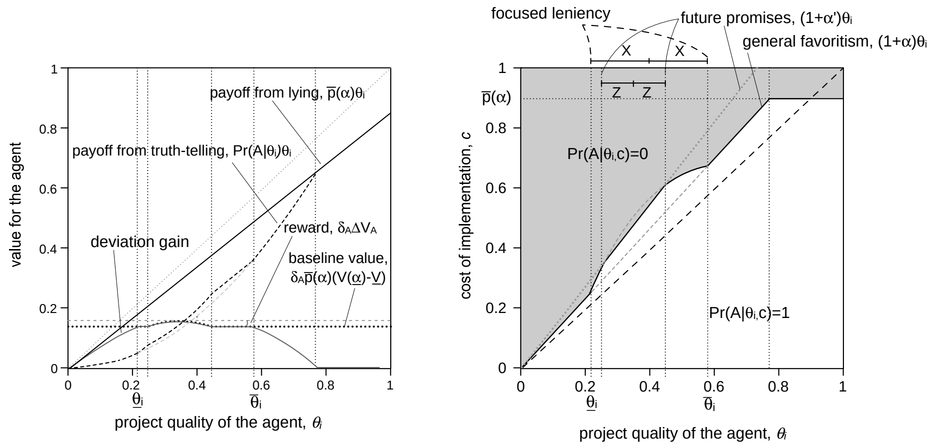

(i) Agent's payoff structure under the second-best decision rule

(ii) Second-best decision rule

Figure 9: Example of a second-best decision rule: $\alpha = 0.1$ and $\overline{\overline{\theta}}_i = 0.8$.

Proof. See Appendix B.4.3 ■

Intuitively, given the assumed two-state structure, the marginal cost and benefit of using the promise of future influence is constant, given by $\delta_P \Delta V_P$ and $\delta_A \Delta V_A$, respectively. In contrast, while the return in terms of relaxing the truth-telling constraint is also linear in $\Pr(A|\theta_i)$, the cost to the principal is convex in $\Pr(A|\theta_i)$. Thus, it is optimal to first use the cheaper source which is the distortion in the current state, and once that becomes too expensive, it becomes optimal to supplement the immediate settling with a promise of additional future influence to restore the truth-telling constraint.

Of course, in the full setting, the principal can also use a transition downwards, just like before. Further, the transition itself will involve a change in the influence instead of a change in probability of transition, meaning that the two avenues will generally be blended together through the whole region.

To summarize, the principal has four different avenues for managing the agent's truth-telling constraint. First, she can bias the general decision rule in favor of the agent by α, which increases the value of the relationship to the agent. Second, she can discriminate against his best proposals by refusing to implement any project whenever her cost exceeds a given threshold, $p(\alpha)$. Third, she can further bias the decision rule in favor of the agent when the truth-telling constraint is most binding (intermediate quality projects). Fourth, she can use the promise of additional (and reduction of) future influence to supplement the current influence if the decision would be excessively costly to the principal. This adjusted decision rule is illustrated in Figure 9.

B.3 Stationary equilibrium

Solving even the simplified dynamic problem is computationally cumbersome. Thus, just to illustrate the logic of the distortions I will solve for the stationary equilibrium under a principal with commitment power. In this case, the payoffs of both the agent and the principal can be written in closed form as a function of only α and $\bar{p}$, the level of general favoritism and the maximal acceptance probability.[8] The resulting continuation value then pins down the extent of focused leniency that is needed to maintain truth-telling by the agent.

The solution is illustrated in Figure 10. The key is panel (ii), which illustrates how the distortions grow as the agent becomes increasingly impatient, where both the range $\left[\underline{\theta}, \overline{\theta}\right]$ over which the principal chooses to exercise focused leniency and the range $\left[\overline{\overline{\theta}}, 1\right]$ that the principal discriminates against are growing in the agent's impatience, worsening the equilibrium performance. In short, the increasing impatience of the agent leads the principal to shift the agent's influence from high-quality to average-quality projects. Relatedly, panel (iii) illustrates the extent of general favoritism, where the principal first increases the degree of favoritism to increase the continuation value for the agent, but once the agent becomes sufficiently impatient, the level of general favoritism is decreased because the low value that the agent places on the future makes the higher continuation value an inefficient means of providing value.

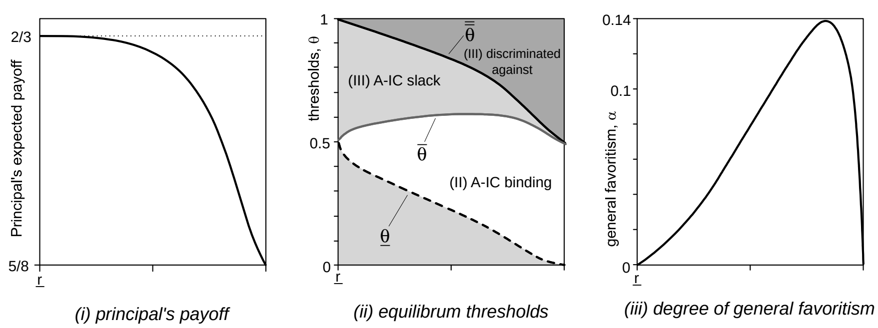

Figure 10: Principal's expected payoff and equilibrium distortions under commitment

The distortions are even more clearly highlighted if we consider how the actual equilibrium decision rule changes as we change the agent's patience. These decision rules are illustrated in Figure 11, which plots the optimal decision rule of the principal for various discount rates of the agent. Intuitively, as the agent initially becomes more impatient, the principal uses all three tools at his disposal. She decreases the maximum acceptance probability to decrease the incentives to exaggerate, thus increasing the discrimination against the best alternatives of the agent. At the same time, to limit the rate at which such discrimination needs to grow, the principal increases the leniency towards the average proposals, increasing the "bulge" in the middle, in relation to the maximum acceptance probability. Finally, the principal initially increases the degree of general favoritism to further increase the agent's continuation value, but eventually decreases it when continuation value

[8] The derivation of these payoffs is not praticularly instructive and available from the author on request.

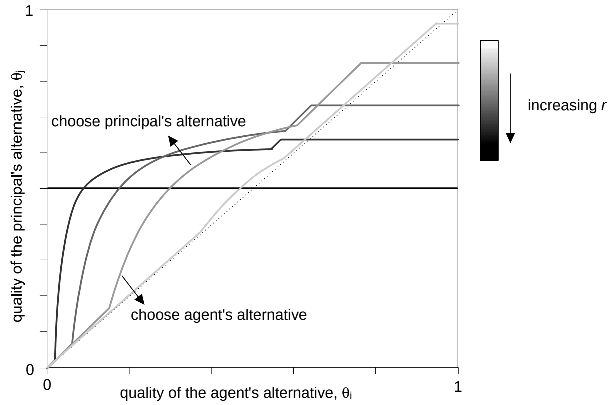

Figure 11: Equilibrium relational influence under principal commitment

becomes a secondary concern to the agent. As the agent becomes infinitely impatient, the decision rule converges to the static optimum: the principal accepts any proposal by the agent whenever her own alternative is worse than average, while choosing her own proposal otherwise.

B.4 Proofs and derivations for the continuous state

B.4.1 Proof of Lemma 3

To find the least-cost means of increasing the continuation value, recall that the agent's expected utility is given by

$$\int_0^1 \theta_i \Pr(A|\theta_i)\, f_A(\theta_i)\, d\theta_i,$$

while the cost of the distortions to the principal (here focusing on favoritism) is given by

$$\int_0^1 \left[\int_0^1 (\theta_i - c)\Pr(A|\theta_i,\theta_P)\, f_P(c)\, dc\right] f_A(\theta_i)\, d\theta_i$$

Two observations follow. First, the minimum cost of delivering any particular $\Pr(A|\theta_i)$ is for the principal to implement the project as long as $c \leq \Pr(A|\theta_i)$ while rejecting the project otherwise. Thus, we can define the expected probability of acceptance by the threshold rule $\overline{c}(\theta_i)$.

Second, on the margin of $\overline{c}(\theta_i)$, consider increasing $\overline{c}(\theta_i)$ while decreasing $\overline{c}(\theta_i')$ in a way that keeps the expected cost to the principal constant. This implies that

$$(\theta_i - \overline{c}(\theta_i))\, f_P(\overline{c}(\theta_i))\, f(\theta_i) + \left(\theta_i' - \overline{c}(\theta_i')\right) f(\theta_i')\, f_P\left(\overline{c}(\theta_i')\right) \frac{d\overline{c}(\theta_i')}{d\overline{c}(\theta_i)} = 0 \Leftrightarrow \frac{d\overline{c}(\theta_i')}{d\overline{c}(\theta_i)} = -\frac{(\theta_i - \overline{c}(\theta_i)) f(\theta_i) f_P(\overline{c}(\theta_i))}{\left(\theta_i' - \overline{c}(\theta_i')\right) f(\theta_i') f_P\left(\overline{c}(\theta_i')\right)}.$$

The corresponding effect on the agent's expected payoff is given by

$$\int_0^1 \left[\int_0^{\overline{c}(\theta_i)} \theta_i f_P(c)\, dc\right] f_A(\theta_i)\, d\theta_i$$

$\theta_i f_P\left(\overline{c}\left(\theta_i\right)\right) f\left(\theta_i\right)+\theta_i' f_P\left(\overline{c}\left(\theta_i'\right)\right) f\left(\theta_i'\right) \frac{d\overline{c}\left(\theta_i'\right)}{d\overline{c}\left(\theta_i\right)}$.

Now, holding the principal's payoff constant, $\frac{d\overline{c}\left(\theta_i'\right)}{d\overline{c}\left(\theta_i\right)}=-\frac{\left(\theta_i-\overline{c}\left(\theta_i\right)\right) f\left(\theta_i\right) f_P\left(\overline{c}\left(\theta_i\right)\right)}{\left(\theta_i'-\overline{c}\left(\theta_i'\right)\right) f\left(\theta_i'\right) f_P\left(\overline{c}\left(\theta_i'\right)\right)}$, and thus optimality of the policy requires that

$$\theta_i f_P\left(\overline{c}\left(\theta_i\right)\right) f\left(\theta_i\right)-\theta_i' f_P\left(\overline{c}\left(\theta_i'\right)\right) f\left(\theta_i'\right) \frac{\left(\theta_i-\overline{c}\left(\theta_i\right)\right) f\left(\theta_i\right) f_P\left(\overline{c}\left(\theta_i\right)\right)}{\left(\theta_i'-\overline{c}\left(\theta_i'\right)\right) f\left(\theta_i'\right) f_P\left(\overline{c}\left(\theta_i'\right)\right)}=0.$$

This then rearranges then to

$$\frac{\theta_i}{\theta_i'}=\frac{\overline{c}\left(\theta_i\right)}{\overline{c}\left(\theta_i'\right)} \Rightarrow \overline{c}\left(\theta_i\right)=(1+\alpha)\, \theta_i, \alpha \geq 0. \tag{46}$$

B.4.2 Leniency towards average proposals

Given the level of general favoritism, discrimination against the best proposals, and continuation value, we can write the agent's truth-telling constraint as

$$\delta_A \Delta V^{cont} \geq\left(1-\frac{(1+\alpha)\, \theta_i}{\overline{p}\,(\alpha)}\right) \theta_i, \tag{47}$$

where ΔV^{cont} is the expected loss from lying. Next, using the concavity of the deviation temptation, maximized at $\widehat{\theta}_i=\frac{\overline{p}(\alpha)}{2(1+\alpha)}$, we can observe that if the constraint is violated at $\widehat{\theta}_i$, then there exist $\underline{\theta}<\widehat{\theta}_i<\overline{\theta}$ such that

$$\delta \Delta V^{cont}=\left(1-\frac{(1+\alpha) \underline{\theta}}{\overline{p}(\alpha)}\right) \underline{\theta}=\left(1-\frac{(1+\alpha) \overline{\theta}}{\overline{p}(\alpha)}\right) \overline{\theta}.$$

In other words, while for $\theta_i \leq \underline{\theta}$, the project is sufficiently worthless that exaggeration is not worth the cost of destroyed reputation, for $\theta_i \geq \overline{\theta}$, the project is sufficiently good so that the incremental increase in the acceptance probability is not worth destroying the reputation.[9] To satisfy the truth-telling constraint for higher types, it then needs to be that

$$\left(1-\frac{\operatorname{Pr}\left(A \mid \theta_i\right)}{\overline{p}(\alpha)}\right) \theta_i=\left(1-\frac{(1+\alpha) \underline{\theta}_i}{\overline{p}(\alpha)}\right) \underline{\theta},$$

which allows us to write $\operatorname{Pr}\left(A \mid \theta_i\right)$ as

$$\operatorname{Pr}\left(A \mid \theta_i\right)=\overline{p}\,(\alpha)-(\overline{p}\,(\alpha)-(1+\alpha)\, \underline{\theta}) \frac{\underline{\theta}}{\theta_i}. \tag{48}$$

Further, the thresholds themselves are given by

[9] Without further analysis, it appears possible that $\overline{\theta}_i<\overline{\overline{\theta}}_i$, so that the upper threshold doesn't exist. Below, I show that the structure of the problem implies that $\overline{\theta}_i<\overline{\overline{\theta}}_i$.

$$\{\underline{\theta}, \overline{\theta}\} = \frac{\overline{p}(\alpha)}{2(1+\alpha)}\left(1 \pm X\right), \qquad X = \sqrt{1 - \frac{4(1+\alpha)\delta_A \Delta V^{cont}}{\overline{p}(\alpha)^2}}.$$

B.4.3 Proof of Lemma 4

From the agent's truth-telling constraint, we get that the tradeoff between the acceptance probability and the transition probability needs to satisfy

$$\frac{d\Pr(A|m_i,\underline{\alpha})}{dq(\theta_i)} = -\frac{\delta_A \Delta V_A}{\theta_i}$$

For the principal, the first-order condition for the optimal choice is given by

$$\Delta\delta_P \widetilde{V}_P + \frac{d\Pr(A|m_i,\underline{\alpha})}{dq(\theta_i)}\theta_i - \Pr\left(A|m_i,\underline{\alpha}\right)\frac{d\Pr(A|m_i,\underline{\alpha})}{dq(\theta_i)} = 0,$$

or $\frac{\Delta\delta\widetilde{V}_P}{(\Pr(A|m_i,\underline{\alpha})-\theta_i)} = \frac{d\Pr(A|m_i,\underline{\alpha})}{dq(\theta_i)}$. Noting that $\Delta\widetilde{V}_P < 0$, define $\Delta V_P = -\Delta\widetilde{V}_P$ as the cost of transitioning to the better state for the principal. Then we have that, since both derivatives need to be satisfied (otherwise we could readjust to improve performance), we get that the optimal point at which we use the future value to compensate for the current value is

$$-\frac{\delta_A \Delta V_A}{\theta_i} = -\frac{\delta_P \Delta V_P}{(\Pr(A|m_i,\underline{\alpha})-\theta_i)} \Leftrightarrow \Pr\left(A|m_i,\underline{\alpha}\right) = \left[1 + \frac{\delta_P \Delta V_P}{\delta_A \Delta V_A}\right]\theta_i.$$

Intuitively, $\frac{\delta_P \Delta V_P}{\delta_A \Delta V_A}$ is the cost-benefit ration of considering a particular continuation equilibrium. As long as the distortion is less costly, use the current continuation value, as the distortion needed grows, better to utilize the promise of future continuation value.