BIROn - Birkbeck Institutional Research Online

Yu, Tinghua (2021) Intrinsic motivation, office incentives, and innovation. Working Paper. Birkbeck Centre for Applied Macroeconomics, London, UK.

Downloaded from: https://eprints.bbk.ac.uk/id/eprint/43390/

Usage Guidelines:
Please refer to usage guidelines at https://eprints.bbk.ac.uk/policies.html or alternatively contact lib-eprints@bbk.ac.uk.

ISSN 1745-8587

Birkbeck Centre for Applied Macroeconomics

BCAM 2106

Intrinsic Motivation, Office Incentives, and Innovation

Tinghua Yu
Birkbeck, University of London

March 2021

▪ Birkbeck, University of London ▪ Malet Street ▪ London ▪ WC1E 7HX ▪

Intrinsic Motivation, Office Incentives, and Innovation

Tinghua Yu *

Abstract

Many organizations, such as government agencies and NGOs, learn about policy effectiveness through decentralized experimentation. However, unobserved effort by an agent can affect the outcome of an experiment, thus limiting its informativeness. A principal can improve the informativeness of an experiment by motivating the agent, using office as an incentive. I develop a model of office incentives in a decentralized experimentation setting where agents are motivated by organizational goals. The principal may keep the agent in office only when the outcome of an experiment is good, thereby creating high-powered office incentives for the agent. High-powered office incentives motivate the agent's effort in implementing the experiment in order to stay in office. However, they also reduce the agent's expected informational benefits from experimentation, which can reduce the effort expended by the agent in implementing the experiment. The degree to which the agent values achieving organizational goals affects such trade-offs. I show that the principal is more likely to use high-powered incentives when the agent places a high value on achieving organizational goals and when multiple agents implement the same experiment.

*Email address: t.yu5@lse.ac.uk; Department of Government, London School of Economics and Political Science. I thank Alessandra Casella, Tim Frye, Shigeo Hirano, Alex Hirsch, John Huber, Macartan Humphreys, Navin Kartik, Kimuli Kasara, Yotam Margalit, Pablo Montagnes, Massimo Morelli, Suresh Naidu, Andy Nathan, Salvatore Nunnari, Pietro Ortoleva, Carlo Prato, Richard Van Weelden, Stephane Wolton, Mike Ting, seminar participants at Columbia University Comparative Politics Seminar and Political Economy Breakfast Seminar, and panelists at the 2015 Annual APSA Conference, the 2015 Annual MPSA Conference, and the 2015 Annual SPSA Conference. All remaining errors are my own.

Many organizations, such as government agencies and NGOs, learn through experimentation. Outcomes of experimentation often depend significantly on unobserved effort decisions made by agents. For instance, if a legislature wants to learn about the efficacy of a new education program, individual schools decide how much effort to exert in implementing the experimental program. Similarly, when donors try out a new developmental project, NGOs that implement the project choose their level of effort in it. In China, experimentation has decisively shaped the making of policies in many domains, such as economic reform, inter-party democracy, public education, etc.[1] To learn how effective a new policy is, the authority in the central government often experiments with it at the local level.[2] Local officials implement experimental policies, and the outcomes of experiments provide feedback for the central authority's future policymaking. An ineffective policy produces bad outcomes. An effective policy could also produce bad outcomes if agents shirk in implementation. With the low effort of agents, little could be learned about the efficacy of policy. In such a process, agents' unobserved effort limits the informativeness of an experiment.[3]

To maximize information, principals often use the office itself to motivate the agent to put effort into an experiment. Because public bureaucracies and NGOs do not offer much formal bonus pay for performance, using the office as an incentive is crucial to these organizations. An important component of office value comes from the agent's preference

[1] See Cao, Qian, & Weingast (1999); Fewsmith (2013); Heilmann (2008); Wang (2009); Xu (2011) for discussion on policy experimentation on various issues.

[2] Local-initiated policy experimentation, and center-sponsored experimentation, distinguished by the source of the policy decision, are the two main types of experimentation at the local level. In a local-initiated policy experiment, the local officials make the policy decision to experiment. In center-sponsored experimentation, the central authority imposes experimental policies on the local agents. In both types of experimentation, local officials are responsible for implementation. Scholars debate about whether a specific local experiment is local-initiated or center-sponsored (See Cai & Treisman (2006) for more discussion). The model in the paper helps to understand center-sponsored experimentation.

[3] Hirsch (2016) and Chassang, Miquel, & Snowberg (2012) also discuss the implication of agent's effort decision for learning in experiments.

for achieving organizational goals. The agent may intrinsically share organizational goals.[4] In the public sector, public service motivation is the major source of intrinsic motivation (Francois, 2000; Le Grand, 2006; Perry & Hondeghem, 2008). The agent may also identify with organizational goals other than serving the public good (Akerlof & Kranton, 2005; Besley & Ghatak, 2005; Sheehan, 1996; Wilson, 1989). In addition, whether an organization achieves its goal affects its funding and survival. Downs (1967) argues that "No bureau can survive unless it is continually able to demonstrate that its services are worthwhile to some group with influence over sufficient resources to keep it alive." Because the agent's material well-being hinges on the organization's funding and his career on the organization's survival, the agent is concerned with achieving organizational goals.

How should the principal use office to motivate the agent in decentralized experimentation? Should she keep the agent in office only when the outcome of an experiment is good, thereby creating high-power office incentives or should she keep the agent in office regardless of the outcome? To address these questions, I develop a formal model to analyze the principal's decision of whether to introduce high-powered office incentives in decentralized experimentation. The principal cares about achieving the organizational goal. When in office, the agent also has a preference for achieving the organizational goal. A status quo policy and an experimental policy are available. The effectiveness of the status quo policy in achieving the organizational goal is known. The effectiveness of the experimental policy in achieving the organizational goal is unknown *ex ante*. An effective policy is more likely to achieve the organizational goal if the agent works harder. An ineffective policy always fails. To learn about the effectiveness of the experimental policy, the principal chooses the experimental policy for the agent to implement in the beginning.

[4] The literature conceptualizes intrinsic motivation in two ways. Some consider that individuals obtain payoffs only when they are working on the provision of the policy (Besley & Ghatak, 2005; Andreoni, 1990). Others regard intrinsic motivation as a sort of pure altruistic concern that causes individuals to care about the policy regardless of the policy provider's identity (Francois, 2000; Gailmard & Patty, 2007). I take the first approach in this paper.

The game begins with the principal's decision of whether to retain the agent only when the outcome of the experiment is good or to retain the agent unconditionally. The former type of re-appointment rule creates high-powered office incentives and the latter, low-powered office incentives. The agent sets a level of effort in implementing the experimental policy. At the end of the first period, the policy outcome is revealed to all players. According to the re-appointment rule, the principal retains an agent in office or replaces him with a new agent who shares the preference in achieving the organizational goal with the sitting agent. In the second period, learning from the policy experiment, the principal decides whether to adopt the experimental policy or the status quo policy. An agent in office decides how much effort to expend in implementing the second-period policy.

One building block of the model is that effort expended by the agent in experimentation in the first period affects information about the experimental policy, which is used for policymaking in the second period. A higher level of effort provides better information about the experimental policy. Based on better information, the principal can make better policy decisions in achieving organizational goals, which benefits both the principal and the agent in future office. In other words, both the principal and the agent in future office derive informational benefits from experimentation.

To gain better information, the principal chooses office incentives that motivate the agent's effort in experimentation. On the one hand, if the principal adopts high-powered incentives, the agent may not stay in office. Yet the agent's effort in experimentation leads to his informational benefit in the second period only if he stays in office. Higher-power incentives thus make the agent hold back experimentation effort in the first place. On the other hand, high-powered incentives may also motivate the agent. To stay in office, the agent puts effort into experimentation. When the expected payoff of future office is higher, this motivation effect of high-powered incentives is stronger.

The degree to which the agent values achieving organizational goals affects the principal's trade-offs. When the agent places a low value on achieving organizational goals, the agent is less motivated to exert effort. With low effort, his chance of staying in office to reap learning benefits is low. Thus, he is more likely to hold back effort in experimentation given high-powered incentives. In addition, a lower effort in experimentation leads to a lower informational benefit in future office and hence a lower total expected payoff in future office. The motivation effect of high-powered incentives is weaker. Therefore, high-powered incentives are more likely to dampen incentives for the agent who places a low value on achieving organizational goals. Consequently, the principal refrains from using high-powered incentives when the agent places a low value on achieving organizational goals.

Decentralized experimentation often involves multiple agents. Take China, for an example, where most policy experiments are implemented by local officials in different localities.[5] Likewise, American bureaucracies are replete with of examples in which different agencies or branches within an agency carry out the same task.[6] I extend the basic model to incorporate a situation with two agents. I contrast the principal's choice of incentive structures in a one-agent setting with that in a two-agent setting. In the one-agent environment, only one agent's effort matters for policy learning. In the two-agent environment, both agents' effort contributes to policy learning. If the other agent exerts more effort, an agent's own effort becomes less crucial for policy learning, and the marginal informational benefits of one's own effort diminish. Balancing the cost of effort and its return at the margin, an agent is less concerned with not reaping the learning benefit in this case. In addition, if the other agent exerts more effort, information about an experimental policy improves and an agent's future office becomes more valuable. This strengthens the motivation effect of high-powered incentives. Generally speaking, compared to the one-agent setting, the principal is more

[5] See footnote 1 and footnote 2.

[6] For example, Bendor (1985) discusses issues in welfare policy in the 1960s.

likely to introduce high-powered incentives in the two-agent setting.

This paper contributes to the literature on incentive issues outside a standard private-sector context (Acemoglu, Kremer, & Mian, 2008; Alesina & Tabellini, 2007; Akerlof & Kranton, 2005; Besley & Ghatak, 2005; Benabou & Tirole, 2003; Dixit, 2002; Maskin & Tirole, 2004). It relates to the policy experimentation literature, specifically the strand of literature that examines the aspect of career risk involved with policy innovation (Cai & Treisman, 2009; Cheng & Li, 2015; Majumdar & Mukand, 2004; Rose-Ackerman, 1980). The difference between this paper and that strand of literature is that this paper emphasizes how unobserved effort affects the experimental outcome. In addition, this paper speaks to the literature on policy experimentation in federal systems. One focus of the literature is on the informational externality associated with policy experimentation across regions (Strumpf, 2002; Volden, Ting, & Carpenter, 2008; Callander & Harstad, 2015). Finally, vast literature discusses learning in the private sector (Keller, Rady, & Cripps, 2005; Bolton & Harris, 1999). Some research examines learning in a principal-agent setting where agents are motivated by monetary incentives (Bergemann & Hege, 2005; Bonatti & Hörner, 2011; Halac, Kartik, & Liu, 2016; Manso, 2011).

Model

Environment and Players

The game takes place over two periods, denoted by $t = 1, 2$. A principal P makes a policy choice. Agent A_1 and agent A_2 implement the policies in their jurisdictions. P commits to a re-appointment rule which is based on A_i's policy outcome in period 1. There are two policy options: a status quo policy, denoted by 0, and an experimental policy, denoted by 1. To learn about the experimental policy, P chooses the experimental policy in the first period. A_i sets an effort level $a_{i1} \in [0,1]$ in policy experimentation. In period 2, P chooses a policy $p_2 \in \{0,1\}$. If A_i is re-appointed, he sets an effort level $a_{i2} \in [0,1]$ in implementing policy in period 2. Otherwise, a new appointee implements the policy.

Let x_{it} be the policy outcome in jurisdiction i at period t. If the policy is successful, it yields an outcome of 1. If it fails, it yields an outcome of 0. The effectiveness of policy and effort in implementation jointly determine the policy outcome. The effectiveness of an experimental policy, denoted by θ, is *ex ante* unknown. It could be $\underline{\theta}$ or $\overline{\theta}$. Throughout, I refer to type $\underline{\theta}$ as "ineffective" and type $\overline{\theta}$ as "effective". All players share common prior beliefs about θ, where $\Pr(\theta = \overline{\theta}) = \frac{1}{2}$. If $\theta = \underline{\theta}$, the experimental policy fails. If $\theta = \overline{\theta}$, with probability a_{it}, the policy succeeds; with probability $1 - a_{it}$, the policy fails. When the status quo policy is implemented, the probability of success is γa_{it}. γ thus measures the effectiveness of the status quo policy. It is known to all players that $\gamma \in [\frac{1}{2}, \frac{2}{3}]$.[7] At the end of period t, policy outcome x_{it} is revealed to all players. I summarize the policy outcome as follows.

If $p_t = 0$, the policy outcome x_{it} is distributed as follows.

$$x_{it} = \begin{cases} 1 & \text{with probability } \gamma a_{it} \\ 0 & \text{with probability } 1 - \gamma a_{it}. \end{cases} \tag{1}$$

[7] When $\gamma \geq \frac{1}{2}$, the *ex ante* outcome of the experimental policy in the first period is not better than the outcome of the status quo policy. When $\gamma > \frac{2}{3}$, even if the experimental policy is revealed to be effective *ex post*, the principal is better off adopting the status quo in the first period. Thus, the lower bound ensures that the principal undertakes experimentation in the first period to learn about the experimental policy, and the upper bound ensures that experimenting in the first period is possibly beneficial for the principal in the long run.

If $p_t = 1$, the policy outcome x_{it} is distributed as follows.

$$x_{it} = \begin{cases} 1 & \text{with probability } a_{it} \text{ if } \theta = \overline{\theta}; \text{ with probability 0 if } \theta = \underline{\theta} \\ 0 & \text{with probability } 1 - a_{it} \text{ if } \theta = \overline{\theta}; \text{ with probability 1 if } \theta = \underline{\theta}. \end{cases} \quad (2)$$

In the beginning, P commits to a re-appointment rule that specifies a threshold of the first-period policy outcome, denoted by $\sigma \in \{0, 1\}$. Only if $x_{i1} \geq \sigma$, P re-appoints A_i in the second period. If $\sigma = 0$, P offers low-powered office incentives. If $\sigma = 1$, P provides high-powered office incentives.

P cares about policy outcomes in both jurisdictions, receiving $x_{1t} + x_{2t}$ in period t. P's payoff function is

$$V_P = \sum_i \sum_t x_{it}.$$

A_i cares about policy outcome in his own jurisdiction and receives λx_{it} if he is in office in period t. $\lambda \in [0, 1]$ thus measures the degree to which A_i is motivated by organizational goals. The value of λ is known to all players. A_i incurs a cost of implementation $c(a_{it}) = \frac{a_{it}^2}{2}$. If A_i is replaced, a new agent has the same degree of organizational-goal motivation as A_i. This assumption is to rule out the possibility that P replaces A_i for pure selection reason, and thus to focus on the moral hazard problem. A_i's payoff function is

$$V_{A_i} = \lambda x_{i1} - c(a_{i1}) + I_i(\lambda x_{i2} - c(a_{i2})),$$

where I_i is an indicator function. $I_i = 1$, if A_i stays in office in period 2; and $I_i = 0$, otherwise.

Sequence

This two-period game proceeds as follows.

1. Nature draws the value of θ.

2. P commits to a re-appointment rule σ.

3. A_i chooses a_{i1} in period 1.

4. Nature reveals policy outcomes x_{i1} to P and A_i.

5. P chooses p_2.

6. The agent in jurisdiction i in period 2 chooses a_{i2}.

Solution Concept

This game has a component of information revelation, so I derive perfect Bayesian equilibria in pure strategies. I focus on symmetric equilibrium where both agents adopt the same strategies. Let H^1 be the set of all period 1 histories. The equilibrium consists of strategies: σ, a_{it}, p_2, and beliefs about the experimental policy's type. $\sigma \in \{0, 1\}$. $a_{i1} : \{0, 1\} \to [0, 1]$ maps P's threshold choice onto into A_i's effort choice in period 1. $p_2 : H^1 \to \{0, 1\}$ maps the set of period 1 history to period 2 policy choice. $a_{i2} : H^1 \times \{0, 1\} \to [0, 1]$ maps the set of histories leading to period 2 effort choice to period 2 effort choice in jurisdiction i. For each history, players also have beliefs about the probabilities of the experimental policy's type. All players share the same prior belief, denoted by ρ_0. Let ρ_{1j} be player j's posterior belief by the end of period 1, where $j \in \{P, A_1, A_2\}$.

Results

To show the cost and benefit of high-powered incentives in generating informative experimentation, I begin with an example of one agent. Then, I consider the case of two agents. In each case, I first derive players' strategies in period 2 and describe how effort in experimentation in period 1 affects decisions in period 2. Then, I analyze the agents' strategies in period 1 under different incentive structures. The principal's choice of incentive structures is then discussed. Finally, I compare incentive structures in equilibrium across two cases.

One Agent

In the basic setup, the notations are developed for a two-agent case. Here, I make some necessary notational changes for a one-agent case. An agent is denoted by A, his effort in period t by a_t, policy outcome in period t by x_t, P's second-period choice by p_2^s, and her re-appointment rule by σ^s.

Period 2 Decisions

A key feature of the model is that information available to players in period 2 is endogenous to effort into experimentation in period 1. Suppose that the agent exerts effort a_1 in period 1 in equilibrium. Players update their beliefs over the experimental policy's type using Bayes' rule.[8] If the experiment succeeds, knowing an ineffective policy always fails, all players infer that the experimental policy is an effective type ($\rho_{1j} = 1$). If the experiment fails, it could be caused by an ineffective policy or by insufficient effort. More specifically, the posterior belief that the experimental policy fails despite being effective for player j is

$$\rho_{1j} = \frac{\frac{1}{2}(1-a_1)}{\frac{1}{2}(1-a_1)+\frac{1}{2}} \leq \frac{1}{2}.$$

In the case of experimentation failure, the posterior beliefs of all players are less than or equal to their priors.

Based on the information about the experimental policy, players make second-period

[8] In the appendix, I show that given any belief that the principal could hold off-equilibrium, the agent has no incentive to deviate from his equilibrium action.

decisions. The second-period decisions include the principal's policy choice p_2 and an effort decision a_2 by an agent in office. Because the second period is the last period, the principal chooses a policy that gives her a higher expected payoff in a single period, and the agent exerts effort to maximize his single period payoff. Suppose that the principal adopts the experimental policy in the second period. In this case, an agent in office exerts effort $a_2 = \lambda\rho_1$ and the resulted expected policy payoff is $\rho_1^2\lambda$. If P chooses the status quo policy in the second period, an agent in office exerts effort $a_2 = \lambda\gamma$ and the expected policy payoff in the second period is $\gamma^2\lambda$. If the posterior belief that the experimental policy is an effective type is greater than the effectiveness of the status quo policy ($\rho_1 > \gamma$), the experimental policy yields a higher policy payoff. This condition holds if and only if the first-period experiment succeeds. Observing a successful first-period experiment, the principal adopts the experimental policy in the second period. The following remark summarizes the principal's policy choice in the second period.

Remark 1. *Given the outcome of policy experimentation in period 1, P's period 2 policy choice in the one-agent setting is as follows.*

$$p_2^{s*} = \begin{cases} 1, & \textit{if } x_1 = 1; \\ 0, & \textit{otherwise.} \end{cases}$$

The Learning Premium

As discussed in the previous section, the principal's second-period policy decision depends on information revealed through first-period policy experimentation. Here, I show that policy experimentation is valuable to the principal and the agent who stays in office. The value of policy experimentation depends on how much effort the agent puts into experimentation.

Given the first-period effort a_1 and the prior about the experimental policy, the *ex ante* probability of an experiment being successful is $\frac{1}{2}a_1$. By Remark 1, the *ex ante* probability of the principal adopting the experimental policy in period 2 is $\frac{1}{2}a_1$ and that of choosing the status quo policy is $1 - \frac{1}{2}a_1$. Her expected second-period policy payoff is as follows.[9]

$$E(v(a_1)) = \lambda \frac{1}{2} a_1 (1 - \gamma^2) + \lambda \gamma^2. \tag{3}$$

The term $\lambda \frac{1}{2} a_1 (1 - \gamma^2)$ in the above equation is the principal's learning premium. It represents the effect of the agent's first-period effort in experimentation on the principal's second-period payoff. When an experiment fails, the experimental policy is rejected by the principal in the second period. However, the experimental policy might be effective and the agent's shirking causes the failure. As the agent exerts more effort, the probability that an effective experimental policy is rejected decreases. More effort into experimentation thus increases the learning premium.

If the agent stays in office in the second period, he also benefits from policy experimentation. In addition to the benefit from a better policy decision by the principal in the second period, better information about the experimental policy helps A calibrate his effort better. He works harder for a more effective policy and avoids wasting effort on a less effective policy. Denote the agent's expected payoff in second-period office as $E(w(a_1))$.

$$E(w(a_1)) = \frac{\lambda^2}{2} \frac{1}{2} a_1 (1 - \gamma^2) + \frac{\lambda^2}{2} \gamma^2 \tag{4}$$

As part of the expected payoff in future office, the learning premium of an agent in

[9]The derivation of the following equation is as follows.

$$\begin{aligned} E(v(a_1)) &= \frac{1}{2} a_1 \lambda \rho_{1P}^2 + (1 - \frac{1}{2} a_1) \lambda \gamma^2 \\ &= \frac{1}{2} a_1 \lambda (1 - \gamma^2) + \lambda \gamma^2 \end{aligned}$$

office is captured by the term $\frac{\lambda^2}{2}\frac{1}{2}a_1(1-\gamma^2)$. The higher the degree to which the agent is motivated by organizational goals, the more he values policy experimentation, and the higher leaning premium he receives. The more effective the status quo policy, the less valuable the policy experimentation, and the lower the learning premium.

Low-Powered Office Incentives

First, consider that the principal chooses $\sigma^s = 0$. Regardless of the performance, the agent stays in office and receives the learning premium in the second period. Expecting this, the agent sets an effort level $a_1 \in [0, 1]$ in period 1 to maximize the following objective function.

$$\max_{a_1} \lambda \frac{1}{2} a_1 - c(a_1) + E(w(a_1)).$$

First order condition characterizing the interior solution is as follows.

$$\frac{1}{2}\lambda + \frac{\lambda^2}{2}\frac{1}{2}(1-\gamma^2) = a_1. \tag{5}$$

The right-hand side of the above equation is the marginal cost of effort. The first part in the left-hand side is the current marginal return. The second part in the left-hand side is the marginal learning premium. The agent sets an effort level such that the marginal cost equals the sum of marginal returns in two periods. The following remark summarizes the agent's decision in the first period given low-powered incentives.

Remark 2. *In the subgame where the re-appointment threshold $\sigma^s = 0$, effort in experimentation in period 1 is*

$$a_1^{l*} = \frac{1}{2}\lambda(1 + (1-\gamma^2)\frac{\lambda}{2}),$$

where the superscript denotes that the re-appointment rule provides low-powered office incentives.

Under a re-appointment rule that provides low-powered office incentives, as the agent becomes more motivated by organizational goals, his effort in experimentation in period 1 increases; as the status quo policy becomes more effective, the effort decreases. The current effort leads to better current and future policy outcomes. The stronger the organizational-goal motivation, the more the agent values policy outcomes, and the more effort the agent exerts. When the status quo policy is more effective, learning about the experimental policy becomes less beneficial.

High-Powered Office Incentives

Now, consider that the principal sets $\sigma^s = 1$. She rewards the agent's good performance with future office. The agent's effort in experimentation contributes to good performance. In addition, his effort affects the learning premium which is part of the expected payoff in future office. The following optimization problem characterizes A's effort choice in the first period.

$$\max_{a_1} \lambda \frac{1}{2} a_1 - c(a_1) + \frac{1}{2} a_1 E(w(a_1)).$$

The agent chooses an effort level according to the following first order condition:

$$\frac{1}{2}\lambda + \frac{1}{2}\frac{\lambda^2}{2}(\frac{1}{2}a_1(1-\gamma^2)+\gamma^2) + \frac{1}{2}a_1\frac{\lambda^2}{2}\frac{1}{2}(1-\gamma^2) = a_1. \tag{6}$$

The right-hand side of the above equation is the marginal cost of effort. The marginal current return is captured in the first term in the left-hand side. The future marginal return has two components. The first component, represented in the second term in the left-hand side, is the marginal increase in the probability of staying in office times the expected payoff in future office. With high-powered office incentives, good performance is rewarded with future office. A higher expected payoff in future office provides stronger incentives to work today. The second component, represented in the third term in the left-hand side, is the marginal increase in learning premium, holding the expected probability of staying in office constant. Given high-powered office incentives, the agent also faces uncertainty in reaping the learning premium. The uncertainty plays a larger influence on his effort decision when the marginal learning premium is higher. Balancing the marginal cost and benefit, I derive the agent's equilibrium strategy in the subgame where the threshold $\sigma = 1$ as follows.

Remark 3. *In the subgame where the re-appointment threshold $\sigma^s = 1$, effort in experimentation in period 1 is:*

$$a_1^{h*} = \frac{\frac{1}{2}\lambda(1 + \frac{\lambda}{2}\gamma^2)}{1 - \frac{\lambda^2}{2}\frac{1}{2}(1-\gamma^2)},$$

where superscript denotes that the re-appointment provides high-powered incentives.

With high-powered office incentives, both the organizational-goal motivation and the effectiveness of the status quo policy have positive effects on experimentation effort. Intuitively, an agent who is highly motivated by organizational goals works harder. But why does an agent put more effort into experimentation if the status quo policy becomes more effective? On the one hand, when the status quo policy becomes more effective, the learning premium becomes smaller, and the agent's tendency to shirk in experimentation increases. One the other hand, when the status quo policy becomes more effective, the payoff in future office increases, and the agent tends to work harder to attain office. Because the agent reaps a learning premium with probability $\frac{1}{2}a_1$, his tendency towards shirking is discounted by $\frac{1}{2}a_1$. Overall, the agent works harder when the status quo policy is more effective.

The Principal's Choice of Incentive Structures

As established in Equation (3), the agent's effort in experimentation in the first period contributes to the principal's expected second-period payoff. Moreover, the agent's effort in the first period increases the probability of a good first-period policy outcome and thus the principal's expected first-period payoff. As a result, the principal chooses an incentive structure that induces more effort in experimentation.

A comparison between Equation (5) and Equation (6) demonstrates the cost and benefit of high-powered office incentives. On the one hand, by rewarding good performance with future office, high-powered office incentives motivate the agent to put effort into experimentation. The motivation effect is captured as $\frac{1}{2}\frac{\lambda^2}{2}(\frac{1}{2}a_1(1-\gamma^2)+\gamma^2)$ in Equation (5). The effect is greater when the value of future office is higher. On the other hand, high-powered office incentives introduce uncertainty in reaping the learning premium and thus discourage effort in policy experimentation. Given high-powered office incentives, the agent could only benefit from learning if he stays in office. His expected marginal learning premium is $\frac{1}{2}a_1\frac{\lambda^2}{2}\frac{1}{2}(1-\gamma^2)$. Provided with low-powered office incentives, the agent benefits from learning with certainty and receives a learning premium of $\frac{\lambda^2}{2}\frac{1}{2}(1-\gamma^2)$. The agent's expected marginal learning premium given high-powered office incentives is $(1-\frac{1}{2}a_1)\frac{\lambda^2}{2}\frac{1}{2}(1-\gamma^2)$ less than that given low-powered office incentives. $(1-\frac{1}{2}a_1)\frac{\lambda^2}{2}\frac{1}{2}(1-\gamma^2)$ thus represents the cost of high-powered incentives. As the marginal learning premium becomes greater, the cost becomes larger.

The degree to which the agent values achieving organizational goals affects the principal's trade-offs. When the agent places a low value on achieving organizational goals, the agent is less motivated to exert effort. With low effort, his chance of staying in office to reap learning benefits is low. Thus, he is more likely to hold back effort in experimentation. In other words, the cost of high-powered incentives is larger when the agent is less motivated by organizational goals. In addition, when effort in experimentation is low, the expected payoff in future office is low. The motivation effect is thus small. As a result, the principal refrains from using high-powered incentives when the agent who places a low value on achieving organizational.

To examine formally how the agent's organizational-goal motivation affects the principal's choice of office incentives, I derive the difference between the equilibrium effort under two types of incentive structures as a function of the value the agent places on achieving organizational goals and the effectiveness of the status quo policy.

$$a_1^{h*} - a_1^{l*} = \frac{1}{2}\lambda(\frac{(1+\frac{\lambda}{2}\gamma^2)}{1-\frac{\lambda^2}{2}\frac{1}{2}(1-\gamma^2)} - (1+(1-\gamma^2)\frac{\lambda}{2})))$$

Based on the above expression, I display the overall effect of the agent's value of achieving organizational goals on the relative effectiveness of high-powered office incentives in Figure 1. As the level of the organizational-goal motivation increases, the relative effectiveness of high-powered office incentives first decreases and then increases. As a result, the principal provides low-powered office incentives when the agent places a low value on achieving organizational goals is low and high-powered office incentives when the agent places a high value on achieving organizational. The principal's decision about the incentive structure is stated in the following proposition and illustrated in Figure 2.

Proposition 1.

$$\sigma^{s*} = \begin{cases} 1 & , \textit{ if } \bar{\lambda}^s \leq \lambda \leq 1 \\ 0 & , \textit{ otherwise.} \end{cases}$$

where $\bar{\lambda}^s = \frac{\sqrt{5-8\gamma^2}}{1-\gamma^2} + \frac{1}{\gamma^2-1}$

The principal chooses high-powered office incentives when the agent's value of achieving organizational goals is higher than a threshold, and low-powered incentives otherwise. The threshold is decreasing in the effectiveness of the status quo policy. In other words, the principal is more likely to introduce high-powered office incentives as the status quo policy becomes more effective and as the agent becomes more motivated by organizational goals.

Figure 1: The effect of high–powered incentives with one agent

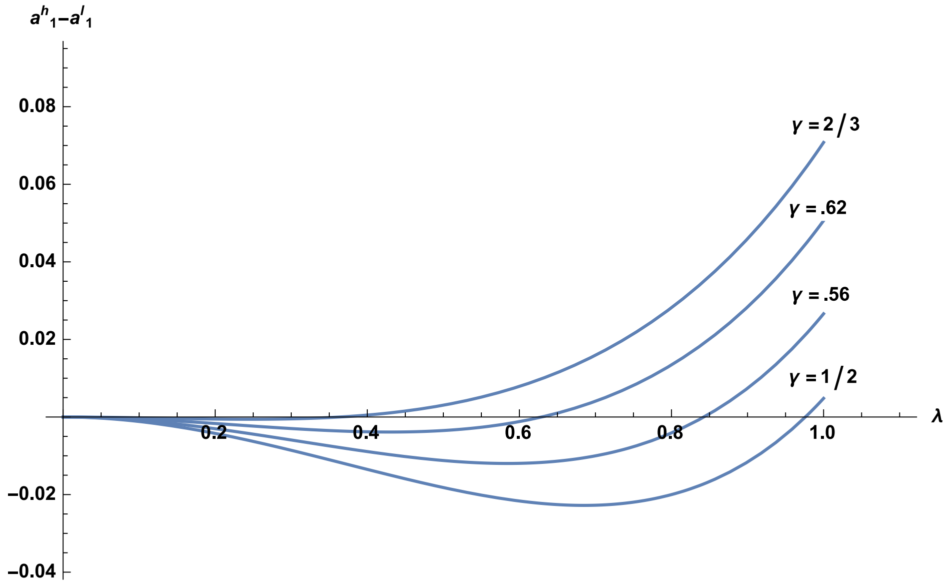

Figure 2: Principal's choice with one agent

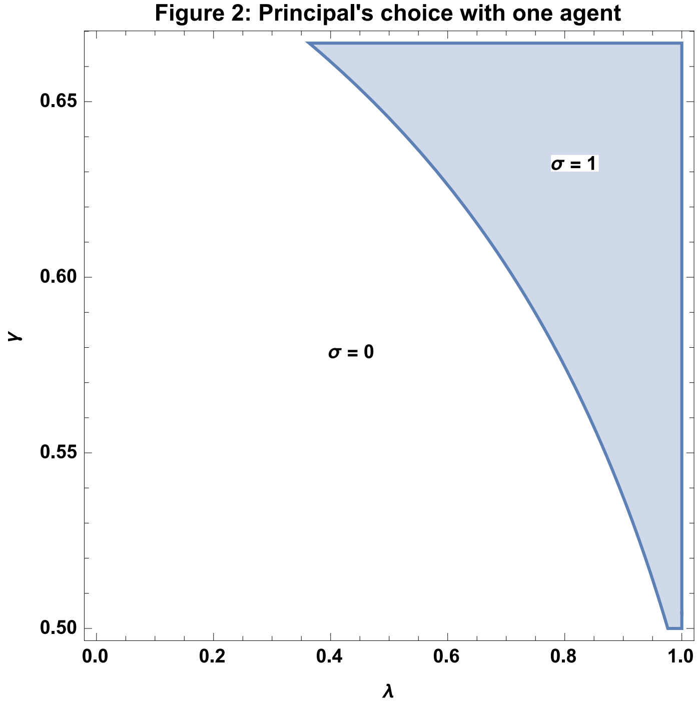

Two Agents

In a two-agent setting, policy learning depends on the agents' joint effort. The other agent's effort also contributes to an agent's future office value. If the other agent exerts more effort, information about an experimental policy becomes better and an agent's future office becomes more valuable. Because of the increase in the expected payoff in future office, the motivation of high-powered office incentives is stronger. At the same time, as the other works harder in experimentation, an agent's own effort becomes less crucial for policy learning, and the marginal learning premium of an agent's effort diminishes. If provided with high-powered office incentives, an agent is less concerned about not reaping the informational premium. The cost of high-powered office incentives becomes weaker. Generally speaking, the existence of the other agent strengthens the benefit of high-powered office incentives and reduces its cost. The principal who would have not chosen high-powered office incentives in a one-agent environment now adopts high-powered office incentives in a two-agent case. I formally demonstrate the above ideas in the following.

I consider symmetric strategies of two agents. It is useful to denote with a subscript $-i$ parameters belonging to the agent that is not A_i. I start the analysis with players' period 2 decisions and a discussion of the learning premium. Then I derive agents' strategies under each incentive structure. Finally, I analyze the principal's choice of incentive structure.

Period 2 Decisions

In the two-agent setting, the beliefs over the experimental policy's type are updated through policy outcomes in both jurisdictions. If $x_{i1} = 1$ or $x_{-i1} = 1$, $\rho_{1j} = 1$. As long as one jurisdiction observes a successful experiment, all players infer that the experimental policy is an effective type. If $x_{i1} = 0$ and $x_{-i1} = 0$, for any player j, the posterior belief that the experimental policy is effective is

$$\rho_{1j} = \frac{\frac{1}{2}(1 - a_{i1})(1 - a_{-i1})}{\frac{1}{2}(1 - a_{i1})(1 - a_{-i1}) + (1 - \frac{1}{2})} \leq \frac{1}{2}.$$

When experimentation in both jurisdictions fails, for each player, the posterior beliefs about the experimental policy are less than or equal to the prior.

Based on information about the experimental policy, the derivation of the agents' effort choices and the principal's policy choice resembles that in the one-agent case. If the status quo policy is implemented, an agent in future office sets an effort level at $\lambda\gamma$, which results in $\gamma^2\lambda$ policy payoff in expectation. If the experimental policy is implemented, an agent in future office exerts $\lambda\rho_{1p}$ level of effort and the expected policy payoff is $\rho_{1p}^2\lambda$. Clearly, the principal adopts the experimental policy if her posterior belief that the experimental policy is effective is greater than the effectiveness of the status quo policy. As long as one district observes a successful experiment, the principal infers that the experimental policy is an effective type. Therefore, the principal adopts the experimental policy in the second period as long as one of the districts succeeds in experimentation, and the status quo policy otherwise.

Remark 4. *Given the outcomes of policy experimentation in period 1, P's period 2 policy choice in two-agent setting is as follows.*

$$p_2^* = \begin{cases} 1, & \text{if } x_{i1} = 1 \text{ or } x_{-i1} = 1; \\ 0, & \text{otherwise.} \end{cases}$$

The Learning Premium

The agents' effort affects information revelation and thus the learning benefit of policy experimentation. Given the first-period effort level profile $\{a_{i1}, a_{-i1}\}$ and the prior about the experimental policy, the *ex ante* probability of an experiment being successful is $\frac{1}{2}(1-(1-a_{i1})(1-a_{-i1}))$. Following Remark 4, the principal adopts the experimental policy in period 2 with an *ex ante* probability of $\frac{1}{2}(1-(1-a_{i1})(1-a_{-i1}))$ and the status quo policy with an *ex ante* probability of $1-\frac{1}{2}(1-(1-a_{i1})(1-a_{-i1}))$. Her expected second-period policy payoff is

$$E(v(a_{i1},a_{-i1}))=\lambda\frac{1}{2}(1-(1-a_{i1})(1-a_{-i1}))(1-\gamma^2)+\lambda\gamma^2, \tag{7}$$

where $\lambda\frac{1}{2}(1-(1-a_{i1})(1-a_{-i1}))(1-\gamma^2)$ is the principal's learning premium.[10] Each agent's effort in experimentation affects information revelation and thus the quality of policy decision. The effect of an agent's effort a_{i1} on learning premium is diminishing in the other agent's effort a_{-i1}. An agent's marginal contribution to better policy making is diminishing in the other's effort. As long as one experiment succeeds, an effective experimental policy is not rejected by the principal. The more effort by the other agent, the more likely the other agent's experiment is successful, and the less important an agent's own success is to the policy making.

If an agent A_i stays in office in the second period, he also benefits from policy learning. The expected payoff in future office to an agent A_i is

$$E(w(a_{i1},a_{-i1}))=\frac{\lambda^2}{2}\frac{1}{2}(1-(1-a_{i1})(1-a_{-i1}))(1-\gamma^2)+\frac{\lambda^2}{2}\gamma^2. \tag{8}$$

$\frac{\lambda^2}{2}\frac{1}{2}(1-(1-a_{i1})(1-a_{-i1}))(1-\gamma^2)$ is A_i's learning premium. As in the one-agent case, better information helps the principal make a better policy decision, which is also in the interest of an agent. Based on better information, an agent can also calibrate effort better in the second period. The other agent's effort into experimentation contributes to policy

[10]The derivation resembles the one in footnote 9.

learning and thus the value of future office; it also reduces the marginal contribution of an agent's effort to the policy learning.

Low-Powered Office Incentives

When provided with low-powered office incentives, A_i stays in office for two periods. The expected payoff in future office $E(w(a_{i1}, a_{-i1}))$ depends on A_i's effort as well as the other agent's effort. Expecting effort a_{-i1} by the other, A_i exerts an effort $a_{i1} \in [0, 1]$ to solve the following maximization problem.

$$\max_{a_{i1}} \lambda \frac{1}{2} a_{i1} - c(a_{i1}) + E(w(a_{i1}, a_{-i1}))$$

The following first-order condition characterizes A_i's best response.

$$\frac{1}{2}\lambda + \frac{\lambda^2}{2}\frac{1}{2}(1-\gamma^2)(1-a_{-i1}) = a_{i1} \tag{9}$$

A_i's first-period effort increases the value of learning premium. The marginal learning premium, represented by $\frac{\lambda^2}{2}\frac{1}{2}(1-\gamma^2)(1-a_{-i1})$ in the first-order condition, is decreasing in the other agent's effort a_{-1i}. The following remark characterizes A's decision in the first period, given low-powered office incentives.

Remark 5. *In the subgame where the re-appointment threshold $\sigma = 0$, effort in experimentation in period 1 in jurisdiction i is*

$$a_{i1}^{l*} = \frac{1}{2}\lambda \frac{1 + \frac{1}{2}\lambda(1-\gamma^2)}{1 + \frac{1}{4}\lambda^2(1-\gamma^2)},$$

where the superscript denotes that the re-appointment rule provides low-powered office incentives.

As in the one-agent setting, the organizational-goal motivation has a positive effect on the level of effort in experiments in the two-agent environment, and the effectiveness of the status quo policy has a negative effect.

High-Powered Office Incentives

With high-powered office incentives, the probability of staying in office depends on an agent's effort a_{i1} in the first period. The value of future office in jurisdiction i depends on both A_i's effort a_{i1} and A_{-i}'s effort a_{-i1}. A_i chooses his first-period effort according to the following optimization problem.

$$\max_{a_{i1}} \lambda \frac{1}{2} a_{i1} - c(a_{i1}) + \frac{1}{2} a_{i1} E(w(a_{i1}, a_{-i1}))$$

A_i's best response is characterized by the following first-order condition.

$$\frac{1}{2}\lambda + \frac{1}{2}\frac{\lambda^2}{2}(\frac{1}{2}(1-(1-a_{i1})(1-a_{-i1}))(1-\gamma^2)+\gamma^2) + \frac{1}{2}a_{i1}\frac{\lambda^2}{2}\frac{1}{2}(1-\gamma^2)(1-a_{-i1}) = a_{i1} \tag{10}$$

When the principal provides high-powered office incentives, good performance is rewarded with future office. Its value is $\frac{\lambda^2}{2}(\frac{1}{2}(1-(1-a_{i1})(1-a_{-i1}))(1-\gamma^2)+\gamma^2)$. A_{-i}'s effort contributes to A_i's future office value and enhances his incentives to work today. With high-powered incentives, A_i is also concerned about not reaping the learning premium. The marginal value of learning is $\frac{\lambda^2}{2}\frac{1}{2}(1-\gamma^2)(1-a_{-i1})$. Because A_i balances the return and cost of his effort at the margin, what matters is the marginal learning premium. When the marginal learning premium is lower, the concern has less influence on an agent's effort decision. The other's effort a_{-i1} diminishes the marginal learning premium and thus attenuates A_i's concern. As a_{-i1} increases, high-powered office incentives become more effective in inducing A_{i1}'s effort. The following remark summarizes the equilibrium strategy in the subgame where the threshold $\sigma = 1$.

Remark 6. *In the subgame where the re-appointment threshold $\sigma = 1$, effort in experimentation in period 1 in jurisdiction i is*

$$a_{i1}^{h*} = \frac{-(1-\frac{3}{8}\lambda^2(1-\gamma^2)) + \sqrt{(1-\frac{3}{8}\lambda^2(1-\gamma^2))^2 + \lambda^2(1-\gamma^2)\frac{1}{2}\lambda(1+\frac{\lambda}{2}\gamma^2)}}{\frac{\lambda^2}{2}(1-\gamma^2)},$$

where the superscript denotes that the re-appointment provides high-powered office incentives.

In the two-agent case, both the organizational-goal motivation and the effectiveness of the status quo policy have positive effects on an agent's effort in experimentation. As demonstrated in the previous section, the same result holds in the one-agent setting.

Principal's Choice of Incentive Structures

Similar to high-powered office incentives in the one-agent case, high-powered office incentives have costs and benefits in motivating the agent in a two-agent case. The office motivation is represented by the term $\frac{1}{2}\frac{\lambda^2}{2}(\frac{1}{2}(1-(1-a_{i1})(1-a_{-i1}))(1-\gamma^2)+\gamma^2)$ in equation(10). Comparing Equation (9) and Equation (10), an agent's marginal learning premium given high-powered office incentives is less than that given low-powered office incentives. The difference capturing the cost is $(1-\frac{1}{2}a_{i1})\frac{\lambda^2}{2}\frac{1}{2}(1-\gamma^2)(1-a_{-i1})$. As the other agent's effort in experimentation increases, an agent's future office value increases, and the office motivation is stronger. Meanwhile, as the other agent puts more effort, an agent's marginal contribution to policy learning diminishes, and the cost is less.

How does the level of the organizational-goal motivation affect the principal's choice of the incentive structure in the two-agent setting? In addition to an agent's own organizational-goal motivation, the organizational-goal motivation of the other agent also affects the cost and benefit of high-powered office incentives. Homogenous agents share the same level of organizational-goal motivation. When the level of the organizational-goal motivation increases, the other agent tends to put more effort as well. More effort by the other increases an agent's expected payoff in future office and reduces his marginal contribution to policy learning. This strengthens the office motivation but weakens the cost of high-powered office incentives. Formally, the effect of the organizational-goal motivation on the relative effectiveness of high-powered office incentives is as follows.

$$a_{i1}^{h*} - a_{i1}^{l*}$$
$$= \frac{-(1-\frac{3}{8}\lambda^2(1-\gamma^2)) + \sqrt{(1-\frac{3}{8}\lambda^2(1-\gamma^2))^2 + \lambda^2(1-\gamma^2)\frac{1}{2}\lambda(1+\frac{\lambda}{2}\gamma^2)}}{\frac{\lambda^2}{2}(1-\gamma^2)} - \frac{1}{2}\lambda(1+(1-\gamma^2)\frac{\lambda}{2}).)$$

Figure 3 shows how the relative effectiveness of high-powered office incentives changes as the level of the organizational-goal motivation changes, given the effectiveness of the status quo policy. Contrasting Figure 1 and Figure 3, I have two observations. First, similar patterns are evident in both settings. The relationship between the relative effectiveness of high-powered office incentives and the level of the organizational-goal motivation is U-shaped. Second, given the same effectiveness of the status quo policy, the turning point at which the organizational-goal motivation starts to reinforce the relative effectiveness of high-powered office incentives is different in two settings. It is at a lower level in the two-agent setting than that in the one-agent setting.

The following proposition characterizes the principal's choice of the re-appointment rule and Figure 4 illustrates her choice.

Proposition 2.

$$\sigma^* = \begin{cases} 1 & , \textit{if } \bar{\lambda} \leq \lambda \leq 1 \\ 0 & , \textit{otherwise} \end{cases}$$

where

$$\bar{\lambda} = -\frac{1}{2(1+\gamma^2)} - \frac{1}{2}\sqrt{\frac{-9+2\gamma^2+16\gamma^4-2\gamma^6-7\gamma^8}{-1+3\gamma^4-3\gamma^8+\gamma^{12}}}$$
$$+\frac{1}{2}\sqrt{\frac{2}{(1+\gamma^2)^2}+\frac{-4+12\gamma^2}{-1+\gamma^4}+\frac{-4-8\gamma^2+12\gamma^4}{1-\gamma^2-\gamma^4+\gamma^6}+\frac{-90-142\gamma^2-56\gamma^4}{(-1+\gamma^2)(1++\gamma^2)^3\sqrt{\frac{-9+2\gamma^2+16\gamma^4-2\gamma^6-7\gamma^8}{-1+3\gamma^4-3\gamma^8+\gamma^{12}}}}}$$

Similar to the results in the one-agent setting, the principal chooses high-powered office incentives only when the organizational-goal motivation is above a threshold. Given the same effectiveness of the status quo policy, the threshold in the two-agent case is lower than that in a one-agent case.

Figure 3: The effect of high–powered incentives with two–agent

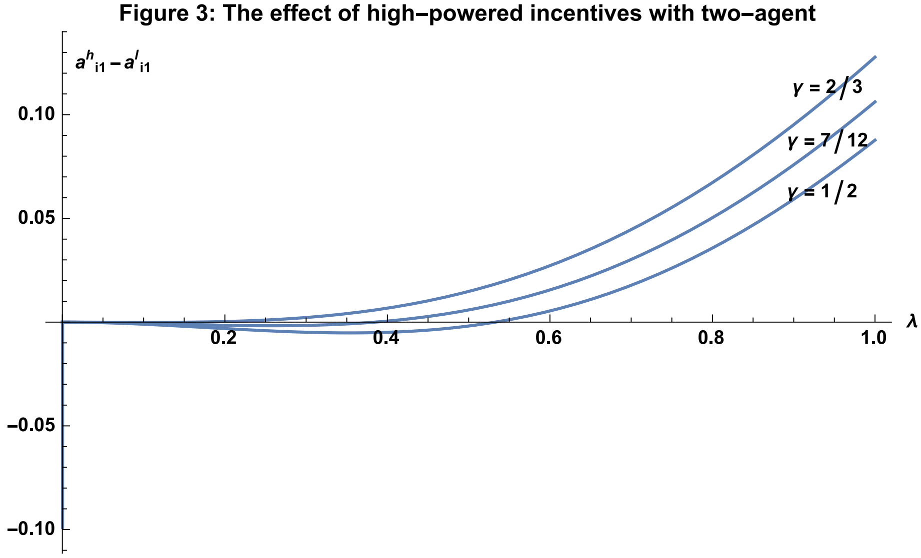

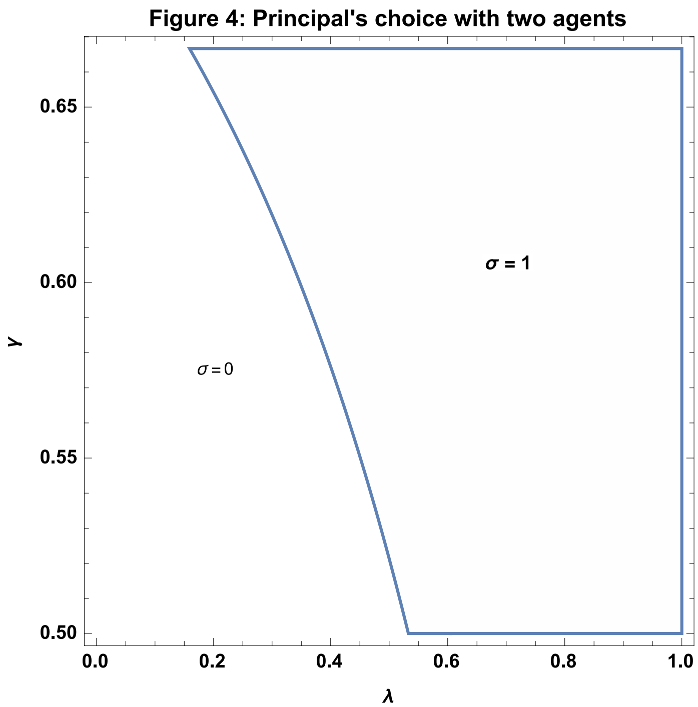

Comparison of Incentive Structures in One-Agent Setting and Two-Agent Setting

Figure 5 contrasts incentive structures in the one-agent setting and that in the two-agent setting. When a status quo policy is effective, high-powered office incentives are introduced in both settings. When a status quo policy is ineffective and policy motivation is low, low-powered office incentives are adopted in both settings. When a status quo policy is ineffective but policy motivation is high, high-powered office incentives are chosen in the two-agent setting and low-powered office incentives in the one-agent setting. Unlike the one-agent setting, high-powered office incentives also induce more effort by the other agent in the two-agent setting, which in turn increases the expected payoff in an agent's future office but reduces an agent's marginal contribution to policy learning. This increases the office motivation and decreases the cost of high-powered office incentives. Therefore, high-powered office incentives are more effective in the two-agent setting. Broadly speaking, the principal is more likely to introduce high-powered office incentives in the two-agent environment.

Figure 5: Incentive structure comparison: one agent vs. two–agent

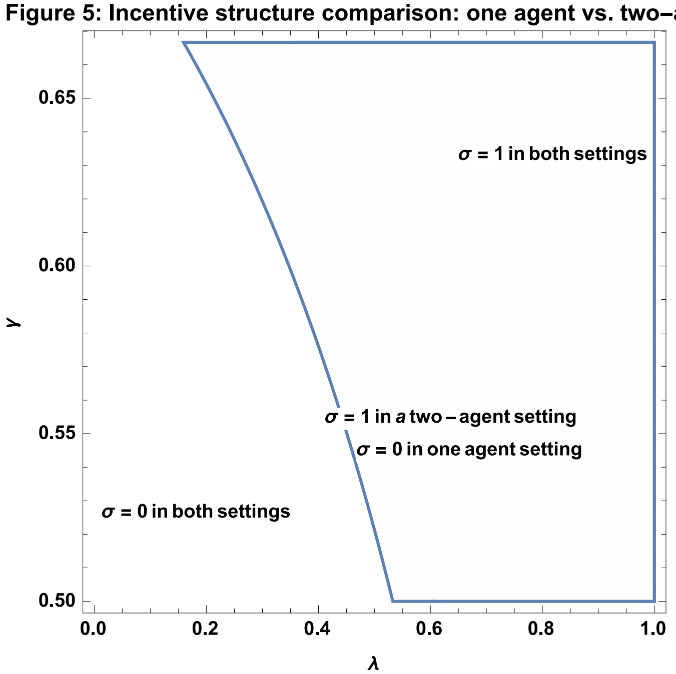

Applications

In this section, I discuss two main contexts in which the model applies. I begin with a discussion of public bureaucracy reform. I then discuss how the implementation of development programs might fit the model.

Public Bureaucracy Reform

The model casts light on issues in the design of incentives in public bureaucracies. By focusing on office as a key incentive and the notion that agents share organizational goals, my approach highlights the incentive issues in public bureaucracies that are different from those in standard private organizations.

The results developed in the model provide some insight into how to offer office incentives when learning is a crucial matter to the organization. Given that agents' intrinsic motivation is a key reason why agents care about achieving organizational goals, this framework implies that office incentives should vary with the degree to which agents are intrinsically motivated. Because of this, the effect of reform of office incentives in public sectors depends on the degree of agents' intrinsic motivation across sectors. In the case of the Chinese bureaucracy, the introduction of high-powered office incentives in environmental agencies and food and drug agencies is an important issue. It is frequently suggested that the policy outcomes are better when the careers of agents in these sectors are tied to the policy outcomes. The model suggests that this type of high-powered incentive is effective provided that these agents care about achieving the organizational goals. However, agents in the Chinese bureaucracy are often rotated across different sectors, and they often do not decide to work in a specific sector. As a result, they might not intrinsically share the goals of a specific organization. The average level of intrinsic motivation in an organization depends on the profile of all agents' career paths. Without considering agents' career paths, attempts to introduce high-powered incentives might reduce the efficiency of these agencies.

Another aspect of public sector reform concerns competition among providers of public goods. As Kaufman (1976) suggests, one of the threats to agencies' survival is competition among agencies for scarce resources. The probability of agency termination is higher when competition among agencies for scarce resources is greatest. A similar logic could hold in the relationship between competition among agencies and agencies' funding. By affecting the survival and funding of an organization, competition changes the degree to which agents care about achieving organizational goals. The approach developed here shows how competition complements high-powered office incentives in the case of learning. High-powered office incentives are more effective when the organization faces fiercer competition. Cross-sectionally, this approach predicts that high-powered office incentives are likely to be negatively correlated with the level of learning when the intensity of competition is low and positively correlated with the level of learning when the intensity of competition is high. It, therefore, seems unsurprising that high-powered office incentives are often used in public organizations that face fierce competition. Within the Chinese bureaucracy, for example, local governments that face intense competition in attracting capital investment often provide high-powered office incentives for agents working on issues of economic development. Under the current Hukou institution which controls population movement, beneficiaries of many public goods are less mobile. This implies a low level of competition among public goods providers across localities. It might explain why low-powered office incentives exist in agencies such as local environmental agencies.

Developmental Programs, NGOs and Government Agencies

It is recognized in the literature that the level of agents' intrinsic motivation varies between NGOs and government agencies. This difference might explain the variation in their performance. As argued by Besley & Ghatak (2001), NGOs may find it easier to screen on motivation than the government and they may also foster public service motivation by providing a better match between the ends of the organization and those of its workers. The electoral concerns of a government imply that some public servants have to carry out policies that they do not believe in. Their public service motivation is thus undermined. Another key issue in the performance variations between the government agencies and NGOs is that of accountability structure. Compared with government agencies, the formal accountability of NGOs is weak. In the context of international development projects, because of the cultural distance between NGOs and local beneficiaries, informal accountability measures, such as social sanctions and enforcement, tends to be weak in the case of NGOs.[11]

When explaining the performance difference between NGOs and government agencies, the existing literature treats the effect of agents' intrinsic motivation and the effect of accountability as separate issues. In terms of the model developed here, different accountability structures correspond to different types of office incentives. This framework thus underlines the complementarity between intrinsic motivation and strong accountability in promoting the performance of developmental programs.

In addition to these empirical implications, the model can also provide some insights into issues in strengthening accountability in both NGOs and government agencies. In a setting where learning is important for social service delivery, the model suggests that strong accountability is suitable for organizations with highly-motivated agents. If it is true that agents in NGOs are on average more motivated than government agents, strengthening accountability in NGO would have a stronger positive effect on performance than strengthening accountability in a government agency. In fact, if government agents are not well motivated intrinsically, strengthening accountability might backfire.

When explaining the performance difference between NGOs and government agencies, the existing literature treats the effect of agents' intrinsic motivation and the effect of the accountability as separate issues. In terms of the model developed here, different accountability structures correspond to different types of office incentives. This framework thus underlines the complementarity between intrinsic motivation and strong accountability in promoting performance of developmental programs.

In addition to these empirical implications, the model can also provide some insights into issues in strengthening accountability in both NGOs and government agencies. In

[11] Social sanctions and enforcement play a decisive role for accountability. Miguel & Gugerty (2005) studies how an inability to impose social sanctions in diverse communities leads to collective action failures in rural western Kenya.

a setting where learning is important for the social service delivery, the model suggests that strong accountability is suitable for organizations with highly-motivated agents. If it is true that agents in NGOs are on average more motivated than government agents, strengthening accountability in NGO would have a stronger positive effect on performance than strengthening accountability in a government agency. In fact, if government agents are not well motivated intrinsically, strengthening accountability might backfire.

Conclusion

The aim of this paper is to explore a principal's decision about agents' office incentives in a setting in which agents' effort in experimentation is crucial for policy learning. The principal constructs office incentives to induce agents' effort in experimentation. High-powered office incentives link good performance to office-holding and thus motivate agents to exert effort. At the same time, by introducing uncertainty in whether the agent could reap the informational benefit from learning, high-powered office incentives also disincentivize effort in experimentation. When agents are highly motivated by organizational goals, the principal is likely to introduce high-powered office incentives. Compared with the one-agent environment, the principal is more likely to provide high-powered office incentives in the two-agent environment.

These ideas are relevant to the discussion of organizations in which agents have a preference for achieving organizational goals. Examples of such organizations include public bureaucracies and NGOs. However, private firms also socialize their employees to share their organizational goals. In future work, it would be valuable to extend this framework to such firms, in order to understand how the interaction between organizational goals and personnel management affects learning and innovation in the private sector.

Appendix

The Mathematica code used for producing graphs is available upon request.

Proof. **Remark 1** Solving Equation (5) gives the result in remark 1. Given the range of λ and γ, the minimum value of a_1^l is 0 and its maximum value is $\frac{11}{16}$. Thus, $a_1^l \in [0, 1]$. □

Proof. **Remark 2** Solving Equation (6) gives the result in remark 2. The minimum value of a_1^h is 0 and its maximum value is $\frac{22}{31}$. Thus, $a_1^h \in [0, 1]$.

□

Proof. **Proposition 1** Solving the following equation system gives the results in proposition 1.

$$\frac{1}{2}\lambda(\frac{(1+\frac{\lambda}{2}\gamma^2)}{1-\frac{\lambda^2}{2}\frac{1}{2}(1-\gamma^2)} - (1+(1-\gamma^2)\frac{\lambda}{2})) \geq 0$$

$$0 < \lambda < 1$$

$$\frac{1}{2} < \gamma < \frac{2}{3}$$

□

Proof. **Remark 3** Solving equation system (7) gives the result in remark 3. The minimum value of a_i^l is 0 and its maximum value is $\frac{11}{19}$. Thus, $a_i^l \in [0, 1]$.

□

Proof. **Remark 4** Solving equation system (8) gives the result in remark 4. The minimum value of a_i^h is 0 and its maximum value is $\frac{1}{20}(-57 + \sqrt{5009})$. Thus, $a_i^h \in [0, 1]$.

□

Proof. **Proposition 2** Solving the following equation system gives the results in proposition

2.

$$\frac{-(1-\frac{3}{8}\lambda^2(1-\gamma^2))+\sqrt{(1-\frac{3}{8}\lambda^2(1-\gamma^2))^2+\lambda^2(1-\gamma^2)\frac{1}{2}\lambda(1+\frac{\lambda}{2}\gamma^2)}}{\frac{\lambda^2}{2}(1-\gamma^2)}-\frac{1}{2}\lambda(1+(1-\gamma^2)\frac{\lambda}{2}))\geq 0$$

$$0<\lambda<1$$

$$\frac{1}{2}<\gamma<\frac{2}{3}$$

□

Proof. **Off Equilibrium Beliefs**

A observes her own action. It is reasonable to suppose that A updates belief about the experimental policy according to Bayes' rule. Now consider P's off-equilibrium beliefs. I prove that to sustain Perfect Bayesian Equilibrium derived in the main section no restriction on P's off-equilibrium beliefs is required. In other words, given any belief that P might hold off-equilibrium, A has no incentive to deviate from his equilibrium action.

Low-Powered Office Incentives

Consider the subgame where low-powered incentives are introduced.

Suppose A deviates to a_1. Following a failure and a success, A forms correct beliefs using Bayes' rule as follows.

$$\rho_{1A}^{f}=\frac{.5(1-a_1)}{.5(1-a_1)+.5}<\gamma$$

$$\rho_{1A}^{s}=1>\gamma$$

If P forms correct beliefs, she chooses $p_2=0$ following a failure and $p_2=1$ following a success. A makes a payoff of $\frac{(\lambda\gamma)^2}{2}$ in period 2 if the experimental fails and a payoff of $\frac{\lambda^2}{2}$ if it succeeds.

Now, consider the off-equilibrium beliefs for P. If a failure is observed, P's off-equilibrium belief is denoted by ρ^f_{1P}, and a success ρ^s_{1P}. To break ties, I assume that P adopts status quo policy if her expected payoff of status quo policy equals to that of the experimental policy. So, P chooses $p_2 = 1$ only if her belief that the probability of the experimental policy being effective is greater than γ. Given P's decision rule in period 2, I classify P's off-equilibrium beliefs into four cases.

1. $\rho^f_{1P} \leq \gamma$ and $\rho^s_{1P} > \gamma$. P makes the same policy decision as she would have if her beliefs are correct. Thus, A receives the same payoff in this case as the payoff he could have received when the principal forms correct beliefs.

2. $\rho^f_{1P} \leq \gamma$ and $\rho^s_{1P} \leq \gamma$. If the experiment fails, P makes the same policy decision as she would have if her beliefs are correct. If the experiment succeeds, given the off-equilibrium belief, P adopts the status quo policy in period 2. A knows that the experimental policy is effective. But his judgement won't matter because the experimental policy won't be adopted. A puts an effort of $\lambda\gamma$ and makes a payoff of $\frac{(\lambda\gamma)^2}{2}$. It is less than what he could have made when the principal forms correct beliefs.

3. $\rho^f_{1P} > \gamma$ and $\rho^s_{1P} > \gamma$. If the experiment succeeds, P makes the same policy decision as she would have if her beliefs are correct. If the experiment fails, given the off-equilibrium belief, P adopts the experimental policy in period 2. A knows that the experimental policy is effective with probability ρ^f_{1A}, so he exerts an effort of $\lambda\rho^f_{1A}$ and makes a payoff of $\frac{(\lambda\rho^f_{1A})^2}{2}$. A thus makes less than than what he could have when the principal forms correct beliefs.

4. $\rho^f_{1P} > \gamma$ and $\rho^s_{1P} \leq \gamma$. P makes the opposite policy decision from what she would have made if her beliefs are correct. It is clear that A makes less than than the payoff he could have received when the principal forms correct belief.

I have shown that for any effort deviating from the equilibrium effort, $a_1 \neq a^l_1$, under any beliefs that P might hold off-equilibrium, A doesn't make a higher payoff than the payoff he receives in the situation where P forms correct beliefs. Thus, I prove that given any off-equilibrium belief of P's, A doesn't receive higher payoff than she would have received in equilibrium.

High-Powered Office Incentives

Now consider the subgame where high-powered office incentives are chosen. Use the same proof strategy as the one in the low-powered office incentives subgame. It could be proved that given any belief that P might hold off-equilibrium A has no incentives to deviate. Two things are worth mentioning. First, if the experiment fails, A makes a payoff of 0 regardless of P's belief. Second, if the experiment succeeds, given A's effort a_1, the probability that A staying office and receiving office value is $\frac{1}{2}a_1$ regardless of P's belief. P's off-equilibrium belief affects A's payoff only by affecting A's office value in period 2, which has been discussed in subgame with low-powered office incentives.

□

References

Acemoglu, Daron, Michael Kremer, & Atif Mian. 2008. "Incentives in Markets, Firms, and Governments." *Journal of Law, Economics, and Organization* 24 (2): 273–306.

Akerlof, George A, & Rachel E Kranton. 2005. "Identity and the Economics of Organizations." *Journal of Economic perspectives* pp. 9–32.

Alesina, Alberto, & Guido Tabellini. 2007. "Bureaucrats or Politicians? Part I: A Single Policy Task." *The American Economic Review* 97 (1): 169–179.

Andreoni, James. 1990. "Impure Altruism and Donations to Public Goods: a Theory of Warm-Glow Giving." *The Economic Journal* pp. 464–477.

Benabou, Roland, & Jean Tirole. 2003. "Intrinsic and Extrinsic Motivation." *The Review of Economic Studies* 70 (3): 489–520.

Bendor, Jonathan B. 1985. *Parallel Systems: Redundancy in Government*. University of California Press.

Bergemann, Dirk, & Ulrich Hege. 2005. "The Financing of Innovation: Learning and Stopping." *RAND Journal of Economics* pp. 719–752.

Besley, Tim, & Maitreesh Ghatak. 2005. "Competition and Incentives with Motivated Agents." *American Economic Review* 95 (3): 616–636.

Besley, Timothy J, & Maitreesh Ghatak. 2001. "Government Versus Private Ownership Of Public Goods." *The Quarterly Journal of Economics* 116 (4): 1343–1372.

Bolton, Patrick, & Christopher Harris. 1999. "Strategic Experimentation." *Econometrica* 67 (2): 349–374.

Bonatti, Alessandro, & Johannes Hörner. 2011. "Collaborating." *The American Economic Review* 101 (2): 632.

Cai, Hongbin, & Daniel Treisman. 2006. "Did Government Decentralization Cause China's Economic Miracle?" *World Politics* 58 (04): 505–535.

Cai, Hongbin, & Daniel Treisman. 2009. "Political Decentralization and Policy Experimentation." *Quarterly Journal of Political Science* 4 (1): 35–58.

Callander, Steven, & Brard Harstad. 2015. "Experimentation in Federal Systems." *The Quarterly Journal of Economics* 130 (2): 951–1002.

Cao, Yuanzheng, Yingyi Qian, & Barry R Weingast. 1999. "From Federalism, Chinese Style to Privatization, Chinese Style." *Economics of Transition* 7 (1): 103–131.

Chassang, Sylvain, Gerard Padro I Miquel, & Erik Snowberg. 2012. "Selective Trials: A Principal-Agent Approach to Randomized Controlled Experiments." *American Economic Review* 102 (4): 1279–1309.

Cheng, Chen, & Christopher Li. 2015. "Laboratories of Democracy: Policy Experimentation under Decentralization.".

Dixit, Avinash. 2002. "Incentives and Organizations in the Public Sector: An Interpretative Review." *Journal of Human Resources* pp. 696–727.

Downs, Anthony. 1967. *Inside Bureaucracy.* Little, Brown Boston.

Fewsmith, Joseph. 2013. *The Logic and Limits of Political Reform in China.* Cambridge University Press.

Francois, Patrick. 2000. "Public Service Motivation as an Argument for Government Provision." *Journal of Public Economics* 78 (3): 275–299.

Gailmard, Sean, & John W Patty. 2007. "Slackers and Zealots: Civil Service, Policy Discretion, and Bureaucratic Expertise." *American Journal of Political Science* 51 (4): 873–889.

Halac, Marina, Navin Kartik, & Qingmin Liu. 2016. "Optimal contracts for experimentation." *The Review of Economic Studies* 83 (3): 1040–1091.

Heilmann, Sebastian. 2008. "Experimentation under Hierarchy: Policy Experiments in the Reorganization of China's State Sector, 1978–2008." *Harvard University Center for International Development Working Paper* 172.

Hirsch, Alexander V. 2016. "Experimentation and Persuasion in Political Organizations." *American Political Science Review* 110 (1): 68–84.

Kaufman, Herbert. 1976. *Are Government Organizations Immortal?* Brookings Institution Press.

Keller, Godfrey, Sven Rady, & Martin Cripps. 2005. "Strategic Experimentation with Exponential Bandits." *Econometrica* 73 (1): 39–68.

Le Grand, Julian. 2006. "Motivation, Agency, and Public Policy: of Knights and Knaves, Pawns and Queens." *OUP Catalogue* .

Majumdar, Sumon, & Sharun W Mukand. 2004. "Policy Gambles." *American Economic Review* pp. 1207–1222.

Manso, Gustavo. 2011. "Motivating Innovation." *The Journal of Finance* 66 (5): 1823–1860.

Maskin, Eric, & Jean Tirole. 2004. "The Politician and the Judge: Accountability in Government." *American Economic Review* pp. 1034–1054.

Miguel, Edward, & Mary Kay Gugerty. 2005. "Ethnic Diversity, Social Sanctions, and Public Goods in Kenya." *Journal of Public Economics* 89 (11): 2325–2368.

Perry, James L, & Annie Hondeghem. 2008. *Motivation in Public Management: The Call of Public Service*. Oxford University Press.

Rose-Ackerman, Susan. 1980. "Risk Taking and Reelection: Does Federalism Promote Innovation?" *The Journal of Legal Studies* pp. 593–616.

Sheehan, Robert M. 1996. "Mission Accomplishment as Philanthropic Organization Effectiveness: Key Findings from the Excellence in Philanthropy Project." *Nonprofit and Voluntary Sector Quarterly* 25 (1): 110–123.

Strumpf, Koleman S. 2002. "Does Government Decentralization Increase Policy Innovation?" *Journal of Public Economic Theory* 4 (2): 207–241.

Vivalt, Eva et al. 2015. "Heterogeneous Treatment Effects in Impact Evaluation." *American Economic Review* 105 (5): 467–70.

Volden, Craig, Michael M Ting, & Daniel P Carpenter. 2008. "A Formal Model of Learning And Policy Diffusion." *American Political Science Review* 102 (03): 319–332.

Wang, Shaoguang. 2009. "Adapting by Learning The Evolution of China's Rural Health Care Financing." *Modern China* 35 (4): 370–404.

Wilson, James Q. 1989. *Bureaucracy: What Government Agencies Do and Why They Do It*. Basic Books.

Xu, Chenggang. 2011. "The Fundamental Institutions of China's Reforms and Development." *Journal of Economic Literature* pp. 1076–1151.