

Yu, Tinghua

Working Paper

Accountability and learning with motivated agents

BCAM Working Paper, No. 2107

Provided in Cooperation with:

Birkbeck Centre for Applied Macroeconomics (BCAM), Birkbeck, University of London

Suggested Citation: Yu, Tinghua (2021) : Accountability and learning with motivated agents, BCAM Working Paper, No. 2107, Birkbeck, University of London, Birkbeck Centre for Applied Macroeconomics (BCAM), London, <https://eprints.bbk.ac.uk/id/eprint/46829>

This Version is available at:

<https://hdl.handle.net/10419/318174>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

BIROn - Birkbeck Institutional Research Online

Yu, Tinghua (2021) Accountability and learning with motivated agents. Working Paper. Birkbeck Centre for Applied Macroeconomics, London, UK.

Downloaded from: <https://eprints.bbk.ac.uk/id/eprint/46829/>

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



BCAM 2107

Accountability and Learning with Motivated Agents

Tinghua Yu

Birkbeck, University of London

September 2021



Accountability and Learning with Motivated Agents

Tinghua Yu *

September 13, 2021

Abstract

Should accountability be introduced to organizations that are learning about the right policies to achieve their goals? I develop an agency model focusing on the interactions between accountability and an agent's intrinsic motivation. More effort by the agent leads to more informative policy outcomes and thereby better policy learning. Holding the agent accountable for the policy outcomes motivates the agent and thus improves policy learning. However, by removing the agent from office upon policy failure and thereby taking away his benefit from learning through failure, accountability also discourages the agent. This negative effect is more substantial when the intrinsic motivation is higher. The principal, therefore, refrains from using accountability on the agent who is more intrinsically motivated.

*Department of Economics, Mathematics and Statistics, Birkbeck, University of London; Email: t.yu@bbk.ac.uk. 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, Ron Smith, 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.

Should a wartime president sack a general who implements an initial strategy but fails to achieve its strategic goals? The conventional wisdom is that accountability serves as an external motivation for better performance, and the president should hold the general accountable. Certainly, the general's implementation of the initial strategy affects the immediate outcome. Yet, in war, the engaging parties often need to learn the right strategy to achieve their goals. Failure of the initial strategy generates valuable information for learning and it is more informative if the general does not slack in implementing the initial strategy. If the general would be sacked upon the failure, the general might be discouraged to work in the beginning because he would not be able to use the knowledge gained from learning through failure to achieve the final success. Accountability, in this case, could backfire.

In the early years of the US civil war, President Lincoln considered whether to sack General George B. McClellan who was making little progress in fighting the war. Eventually, Lincoln held McClellan accountable and dismissed him from his position. McClellan was believed to suffer from a lack of personal motivation. Many organizations face the same question that a wartime president encounters. For example, many government agencies are learning the right policies to govern effectively. Bureaucrats in these agencies implement policies, and the agencies learn through the policy outcomes. Unlike General McClellan, most bureaucrats in developed countries enjoy strong protections from job dismissal. Theory and empirics suggest that bureaucrats are often intrinsically motivated to carry out the task (Besley and Ghatak, 2006; Brehm and Gates, 1999; Prendergast, 2008).

Should an agent implementing policies be held accountable when an organization learns the right policy to achieve its goals? Does accountability crowd in or crowd out intrinsic motivation? The literature has explored how contingent reward (extrinsic motivation) interacting with intrinsic motivation affects effort, but most studies do

not examine settings where experimenting and learning are involved.¹ Among the few exceptions, Gailmard and Patty (2007) argue that job tenure is crucial in inducing intrinsically motivated bureaucrats to make initial investments in policy expertise. Their theory applies to policymakers but not to agents who are implementing policies. This paper aims to fill the gap in the literature and address these questions.

This paper develops a formal model to analyze a principal's decision of whether to introduce accountability on an agent who implements the policy decision when an organization learns the right policy to achieve its goal. The principal (she) cares about achieving the organizational goal. When in office, the agent (he) is intrinsically motivated to achieve the organizational goal and he also receives an additional perk that comes with office. There exists two alternatives. *Ex ante* they are equally likely to be the right policy. The right policy is more likely to achieve the organizational goal if the agent works harder. The wrong policy always fails. The two-period game begins with the principal's decision of whether to hold the agent accountable for the outcome of the initial policy. The agent sets a level of effort in implementing the initial policy. At the end of the first period, the policy outcome is revealed to all players. According to the accountability 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. Learning from implementation of the initial policy, the principal decides the policy for the second period. An agent in office decides how much effort to expend in implementing the second-period policy.

A building block of the model is that the effort expended by the agent on the initial policy contributes to learning. Failure of the initial policy could be caused either by lack of effort or by implementing the wrong policy. Initial effort improves learning through

¹See Besley and Ghatak (2005), Benabou and Tirole (2003), Bénabou and Tirole (2006) and Sliwka (2007) among others.

two mechanisms. First, more initial effort lowers the chance of switching to the wrong policy when the initial policy is, in fact, correct. Second, as the initial effort increases, a failure becomes more informative.

To improve learning, the principal chooses the accountability rule that induces more effort from the agent. Accountability motivates the agent to perform well to stay in office. However, it also discourages the agent. Given accountability, the agent will not remain in office if the initial policy fails. Yet failure provides valuable information, and it is more informative when the initial effort is higher. Knowing that he will not benefit from the learning through failure, the agent holds back the initial effort. A more intrinsically motivated agent tends to exert more effort, which leads to better learning through failure. Therefore, a more intrinsically motivated agent would have more learning benefits to lose if he would be sacked upon the failure of the initial policy. Accountability is more likely to dampen incentives for a more intrinsically motivated agent. As a result, the principal refrains from using accountability on a more intrinsically motivated agent.

The theoretical predictions of the model can inform the design of accountability in many crucial cases. In the discussion section, I discuss the issues of public bureaucracy reform and development programs carried out by NGOs and government agencies. Further, I examine the incentive issues where policy learning involves multiple agents. Each agent implements a policy in his own jurisdiction and policy outcomes from all jurisdictions feed back to the decision by the principal. The existence of another agent does not affect the incentives of an agent who is held accountable for the policy outcome in his jurisdiction but introduces free-riding problems when an agent faces accountability. Generally speaking, compared to the one-agent setting, the principal is more likely to introduce accountability in the two-agent environment.

The paper contributes to the theoretical literature on intrinsic motivation (Ace-

moglu et al., 2008; Alesina and Tabellini, 2007; Akerlof and Kranton, 2005; Besley and Ghatak, 2005; Dixit et al., 2002; Maskin and Tirole, 2004). Benabou and Tirole (2003), Bénabou and Tirole (2006), and Sliwka (2007) analyze how extrinsic motivations crowd out intrinsic motivations in a standard environment without learning. My paper complements the literature by showing how external incentive, accountability more specifically, crowds out intrinsic motivation when the right methods or policies need to be learnt. It shows that by taking away an agent’s benefit of learning through failure, accountability creates misalignment between a principal and an intrinsically motivated agent. This misalignment is greater for a more motivated agent.

The paper also relates to the policy experimentation literature, specifically the strand of literature that examines the aspect of career risk involved with policy innovation (Cai and Treisman, 2009; Cheng and Li, 2019; Majumdar and Mukand, 2004; Rose-Ackerman, 1980). The difference between this paper and that strand of literature is that this paper emphasizes how effort affects the policy outcome. Finally, a vast literature discusses learning in the private sector (Keller et al., 2005; Bolton and Harris, 1999). Some research examines learning in a principal-agent setting where agents are motivated by monetary incentives (Bergemann and Hege, 2005; Bonatti and Hörner, 2011; Halac et al., 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 $p_t \in \{a, b\}$ in each period t . An agent A implements the policy with unobservable effort $e_1 \in [0, 1]$ in period 1. If A is re-appointed, he sets an effort level $e_2 \in [0, 1]$ in implementing policy in period 2. Otherwise, a new appointee implements

the policy in period 2.

If the policy is successful, it yields an outcome of 1. If it fails, it yields an outcome of 0. The policy outcome depends both on whether the policy is right for achieving its goal and an agent’s effort in implementation. To capture this idea, I assume that “nature” initially draws a state $\omega = \{a, b\}$. When the policy matches the state ($p_t = \omega$), the policy is correct; otherwise, it is incorrect. The probability that a correct policy is successful is equal to the implementation effort e_t . The incorrect policy always fails. Players share a common prior belief that $\rho_0 \equiv \Pr(\omega = b) = 1/2$. Without loss of generality, we assume that P adopts $p_1 = a$ in period 1. To break tie, I assume that when P is indifferent between a and b in period 2, she adopts policy b . At the end of period t , policy outcome x_t is revealed to all players.

In the beginning, P commits to whether to hold A accountable for the first-period policy outcome. Denote P ’s decision by $\sigma \in \{0, 1\}$. If $\sigma = 0$, P re-appoints A regardless of first-period policy outcome. If $\sigma = 1$, P re-appoints A only if the first-period policy is successful.

P cares about policy outcomes, receiving x_t in period t . P ’s payoff function is

$$V_P = x_1 + x_2.$$

A intrinsically cares about policy outcome and receives x_t if he is in office in period t . In addition to policy payoffs, A also receives a reward of 1 when in office. This additional office reward could be ego rents or material gains from holding the office. Let A ’s intrinsic motivation be λ , distributed uniformly on $[0, \bar{\lambda}]$.² The value of λ is known to all players. Following Hirsch (2016), I consider that the strength of one’s

² $\bar{\lambda} \equiv \frac{\sqrt{17}-1}{4}$. The assumption ensures that $\bar{\lambda}$ is a necessary and sufficient bound to ensure that if the agent is not held accountable for the first-period policy outcome, the agent’s effort in period 1 is well-behaved and less than 1. It is characterized in the proof of the Proposition.

intrinsic motivation determines one's willingness to put costly effort to achieve good outcome. A incurs a cost of implementation $c(e_t) = \frac{e_t^2}{2\lambda}$. If A is replaced, a new agent has the same degree of intrinsic motivation as A . This assumption is to rule out the possibility that P replaces A for pure selection reason, and thus to focus on the moral hazard problem. A 's payoff function is

$$V_A = x_1 - c(e_1) + 1 + I(x_2 - c(e_2) + 1),$$

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

The sequence of the game is as follows.

Sequence

1. Nature draws the value of ω .
2. P commits to whether to hold A accountable σ .
3. A chooses e_1 in period 1.
4. Nature reveals policy outcomes x_1 to all players.
5. P chooses p_2 .
6. The agent in period 2 chooses e_2 .
7. Nature reveals policy outcomes x_2 to all players.

Solution Concept

This game has a component of information revelation, so I derive perfect Bayesian equilibria in pure strategies. Let H^1 be the set of all period 1 histories. The equilibrium

consists of strategies: σ , e_t , p_2 , and beliefs about the state. $\sigma \in \{0, 1\}$. $e_1 : \{0, 1\} \rightarrow [0, 1]$ maps P 's choice regarding accountability into A 's effort choice in period 1. $p_2 : H^1 \rightarrow \{0, 1\}$ maps the set of period 1 history to period 2 policy choice. $e_2 : H^1 \times \{0, 1\} \rightarrow [0, 1]$ maps the set of histories leading to period 2 effort choice to period 2 effort choice. For each history, players also have beliefs about the state. All players share the same prior belief ρ_0 . To abuse the notation, let ρ be player j 's posterior belief that $\omega = b$ by the end of period 1, where $j \in \{P, A\}$.

Results

I first derive players' strategies in period 2 and describe how effort into the initial policy affects decisions in period 2. Then, I analyze the agent's strategy in period 1 under different accountability structures. The principal's decision over accountability is then discussed.

Period 2 Decision

Players learn about which policy is correct based on policy outcomes in the first period. Using this information, the principal makes a second-period policy decision and the agent in the second period makes an effort decision. Suppose that the agent exerts effort e_1 in period 1 in equilibrium. Players update their beliefs over the underlying state of the world using Bayes' rule.³ If policy a succeeds, knowing that only the right policy can be successful, all players infer that policy a is the right policy ($\Pr(\omega = a) = 1$). In this case, the principal adopts policy a in the second period, and an agent in office exerts effort $e_2 = \lambda$. The policy payoff for the principal is λ , and the net policy payoff for the agent is $\frac{\lambda}{2}$. If policy a fails, it could be caused by an incorrect policy or by

³In 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.

insufficient effort. Formally, the posterior beliefs about the state of the world are

$$\rho(e_1) = \Pr(\omega = b) = \frac{1/2}{1/2(1 - e_1) + 1/2} \in [\frac{1}{2}, 1].$$

The failure of the initial policy a indicates that the alternative policy b is more or equally likely to be the right policy. Notice that $\partial\rho(e_1)/\partial e_1 > 0$. As the initial effort e_1 increases, upon the failure of initial policy a , the posterior belief that policy b is the right policy moves further away from the prior. As the agent puts more effort into the initial policy, the failure of the initial policy becomes more informative. The failure of the initial policy leads to the adoption of the alternative policy b . The agent puts effort $e_2 = \lambda\rho(e_1)$ into the alternative policy b . His effort e_2 is increasing in his prior belief that policy b is the right policy ($\rho(e_1)$). The agent works harder when he is more certain that policy b is the right policy. By improving the informativeness of a failure, the effort into the initial policy leads to more effort in implementing the alternative policy in period 2. The policy payoff for the principal is $\lambda\rho(e_1)$ and the net policy payoff for the agent is $\frac{\lambda\rho^2(e_1)}{2}$.

The Principal Benefits from the Agent's Initial Effort

As discussed in the previous section, based on the learning through the initial policy implementation, the principal makes a policy decision for the second period and the agent the effort decision. Here, I show that the agent's effort into initial policy improves learning and, therefore, benefits the principal.

The principal's second-period expected payoff is:

$$E(v(e_1)) = \frac{1}{2}(e_1\lambda + (1 - e_1)\lambda\rho^2(e_1)) + \frac{1}{2}\lambda\rho^2(e_1).$$

The expression is divided into two terms. The first term, $\frac{1}{2}(e_1\lambda + (1 - e_1)\lambda\rho^2(e_1))$,

is the principal's expected payoff when the initial policy is actually correct and the second term, $\frac{1}{2}\lambda\rho^2(e_1)$, is her expected payoff when the initial policy is wrong. I take the derivative of $E(v(e_1))$ to analyze the two mechanisms through which the principal benefits from the agent's effort into the initial policy.

$$\partial E(v(e_1))/\partial e_1 = \frac{1}{2}\lambda(1 - \rho^2(e_1)) + (1 - \frac{1}{2}e_1)\lambda 2\rho\partial\rho(e_1)/\partial e_1.$$

The first term on the right-hand side of the above equation, $\frac{1}{2}\lambda(1 - \rho^2(e_1))$, is the marginal increase in the principle's second-period policy payoff, holding the expected payoff from policy b constant. The agent's lack of effort into the initial policy increases the chance of failing the right policy and eventually adopting the wrong policy. As the agent's initial effort increases, the likelihood of keeping the initial policy when it is right increases. By improving the principal's policy decision, the agent's initial effort benefits the principal. The second term, $(1 - \frac{1}{2}e_1)\lambda 2\rho\partial\rho(e_1)/\partial e_1$, is the marginal increase in the principle's second period policy payoff from policy b . As discussed previously, upon the failure of policy a , the initial effort improves the learning and thus incentivizes the agent to work harder on implementing policy b . Consequently, conditional on the failure of policy a , the principal's payoff increases in the agent's initial effort.

Without Accountability

First, consider that the principal does not impose accountability on the agent. Regardless of the initial performance, the agent stays in office, receiving the second-period policy payoff and the office reward. Expecting this, the agent sets an effort level $e_1 \in [0, 1]$ in period 1 to maximize the following objective function.

$$\max_{e_1} \frac{1}{2}e_1 - \frac{e_1^2}{2\lambda} + 1 + \frac{1}{2} \left(e_1 \frac{\lambda}{2} + (1 - e_1) \frac{\lambda\rho^2(e_1)}{2} \right) + \frac{1}{2} \frac{\lambda\rho^2(e_1)}{2} + 1.$$

The objective function includes the first-period expected payoff and the second-period expected payoff. The first-period expected payoff includes the first-period expected policy payoff $\frac{1}{2}e_1$, the cost of initial effort $\frac{e_1^2}{2\lambda}$, and the office reward 1. The second-period expected payoff is divided into three terms. The first term, $\frac{1}{2} \left(e_1 \frac{\lambda}{2} + (1 - e_1) \frac{\lambda \rho^2(e_1)}{2} \right)$, is the agent's expected net policy payoff when the initial policy is correct, the second term, $\frac{1}{2} \frac{\lambda \rho^2(e_1)}{2}$, is his expected net policy payoff when the initial policy is wrong, and third term, 1, is the office reward. I examine the first-order condition characterizing the interior solution to analyze the agent's incentives to exert initial effort.

$$\frac{1}{2} + \frac{1}{2} \frac{\lambda}{2} (1 - \rho^2(e_1)) + (1 - \frac{1}{2}e_1) \lambda \rho \partial \rho / \partial e_1 = \frac{e_1}{\lambda}.$$

The right-hand side of the above equation is the marginal cost of the initial effort. The first part on the left-hand side, $\frac{1}{2}$, is the current marginal return. The future marginal return has two components. The first component, $\frac{1}{2} \frac{\lambda}{2} (1 - \rho^2(e_1))$, shows the marginal increase in the agent's second-period net policy payoff, holding the expected net payoff of policy b constant. Conditional on policy a being the right policy, the agent's initial effort increases the likelihood of adopting the right policy in the second period. By improving the principal's policy decision, the agent benefits from his initial effort. The second component, $(1 - \frac{1}{2}e_1) \lambda \rho \partial \rho / \partial e_1$, is the marginal increase in the agent's second period net policy payoff from policy b . Failure of the initial policy a leads to the adoption of the alternative policy b . The initial effort improves the learning through failure and helps the agent calibrate his effort better.

With Accountability

Now, consider that the principal chooses to hold the agent accountable for the failure of the initial policy. The agent stays in office if and only if the initial policy succeeds.

The following optimization problem characterizes A 's effort choice in the first period.

$$\max_{e_1} \frac{1}{2}e_1 - \frac{e_1^2}{2\lambda} + 1 + \frac{1}{2} \left(e_1 \frac{\lambda}{2} + (1 - e_1) \frac{\lambda \rho^2(e_1)}{2} \times 0 \right) + \frac{1}{2} \frac{\lambda \rho^2(e_1)}{2} \times 0 + \frac{1}{2}e_1 \times 1.$$

The objective function includes the first-period expected payoff and the second-period expected payoff. As in the case where an agent faces no accountability, the first-period expected payoff includes the first-period policy payoff $\frac{1}{2}e_1$, the cost of initial effort $\frac{e_1^2}{2\lambda}$, and the office reward 1. To compare an agent's second-period expected payoffs from different accountability structures, I decompose the second-period expected payoff of an agent who is held accountable for the failure of the initial policy into three terms. $\frac{1}{2} \left(e_1 \frac{\lambda}{2} + (1 - e_1) \frac{\lambda \rho^2(e_1)}{2} \times 0 \right)$ is the agent's expected net policy payoff when the initial policy is actually correct. Slacking in the first period causes the failure of the initial policy, which leads to the firing of the agent. Hard working on the initial policy increases the chance of a success, and therefore the chance of staying in office and receiving a payoff of $\frac{\lambda}{2}$. $\frac{1}{2} \frac{\lambda \rho^2(e_1)}{2} \times 0$ is his expected net policy payoff when the initial policy is wrong. No matter how much initial effort the agent puts, the incorrect policy always fails. The agent will be sacked and receives a payoff of 0. $\frac{1}{2}e_1 \times 1$ represents the expected second-period office reward. The agent only receives the office reward when the initial policy succeeds.

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

$$\frac{1}{2} + \frac{1}{2} \frac{\lambda}{2} + \frac{1}{2} \times 1 = \frac{e_1}{\lambda}.$$

The right-hand side of the above equation is the marginal cost of effort. The first part on the left-hand side, $\frac{1}{2}$, is the current marginal return. The future marginal return has two components. The first component, $\frac{1}{2} \frac{\lambda}{2}$, represents the marginal increase

in the agent's second-period net policy payoff. The agent's initial effort decreases the possibility of rejecting the right policy and thereby improves policy decision in the second period. Only when the initial policy succeeds, the agent stays in office and benefits from a better policy decision in the second period. The second component, $\frac{1}{2} \times 1$, is the marginal increase in expected future office reward. Initial effort improves the chance of staying in office and thus receiving office reward.

The Principal's Decision on Accountability

As established, the agent's initial effort contributes to learning and thus the principal's expected second-period payoff. Moreover, by improving policy outcome in the first period, the agent's initial effort increases the principal's expected first-period payoff. As a result, the principal commits to accountability if accountability incentivizes the agent to work harder in period 1.

A comparison between two first-order conditions in the previous sections demonstrates the cost and benefit of accountability. On the one hand, accountability motivates the agent to perform well to stay in office. The motivation effect is captured as $\frac{1}{2}$. On the other hand, accountability could discourage the agent's initial effort. When the agent is held accountable for the initial policy's failure, he only receives future policy payoff when the initial policy succeeds. Yet failure provides useful information. The higher the initial effort is, the more informative failure is. By removing the agent's benefit from learning through failure, accountability could backfire. The agent's expected marginal net policy payoff given accountability is $\frac{1}{(2-e_1)^2} \lambda \frac{3}{4}$ less than that in the alternative case. $\frac{1}{(2-e_1)^2} \lambda \frac{3}{4}$ thus represents the cost of accountability.

The marginal benefit of accountability is the same across different levels of intrinsic motivation. Intrinsic motivation affects the cost $\frac{1}{(2-e_1)^2} \lambda \frac{3}{4}$ both directly and indirectly. Intrinsic motivation λ directly increases the cost. A more intrinsically motivated agent

enjoys higher net policy payoff. By encouraging more initial effort e_1 , intrinsic motivation λ indirectly increases the cost. A more intrinsically motivated agent tends to exert more initial effort, which improves learning through failure. This implies that a more intrinsically motivated agent has more to lose if the initial policy fails and he loses the chance to benefit from learning. Therefore, the disincentives created by accountability are stronger when the intrinsic motivation is higher. As a result, the principal refrains from using accountability on the agent who is more intrinsically motivated. Formally, I summarize the principal's decision in the following proposition.

Proposition. *The principal's decision on accountability is as follows.*

$$\sigma^* = \begin{cases} 1, & \text{if } \lambda \leq .777 \\ 0, & \text{otherwise.} \end{cases}$$

Discussion

In this section, I discuss two main contexts in which the model applies and incentive issues in a two-agent setting.

Public Bureaucracy Reform

The results developed in the model provide some insight into whether to use accountability when learning is a crucial matter to the organization. This framework implies that the use of accountability should vary with the degree to which agents are intrinsically motivated. Because of this, the effect of reform of accountability in public sectors depends on the degree of agents' intrinsic motivation across sectors. In the case of the Chinese bureaucracy, for example, the introduction of accountability in environmental agencies and food and drug agencies is an important issue. It is frequently suggested

that the policy outcomes are better when agents' careers in these sectors are tied to the policy outcomes. The model suggests that accountability is effective provided that agents have a low level of intrinsic motivation. One key feature of the Chinese bureaucracy is that agents are rotated across different sectors and often do not decide which sectors to work. As a result, they might not intrinsically share the goals of a particular organization. The average level of intrinsic motivation in an organization depends on the profile of all agents' career paths. It is lower if more agents have been recently rotated from other posts. This framework suggests that accountability should be introduced in agencies where a sufficient number of agents have been recently rotated.

Developmental Programs, NGOs and Government Agencies

In developing countries, NGOs have been increasingly involved in the provision of relief, welfare, social services, and various development projects.⁴ A growing emphasis on impact evaluation in many of these sectors shows a need for innovation and learning. But the outcomes of similar projects implemented by government agencies and that of NGOs are often different.⁵

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 and Ghatak (2001), NGOs may find it easier to screen on motivation than the government and may also foster public service motivation by providing a better match between the ends of the organization and those of its workers. Because of the electoral concerns of government, some public servants have to carry out policies that they do not believe in. This undermines their

⁴See Besley and Ghatak (2001) for a detailed discussion.

⁵For example, using data on 20 different types of interventions, Vivalt et al. (2015) shows that government-implemented programs also had smaller effect sizes than academic/NGO-implemented programs.

public service motivation. 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, tend to be weak in the case of NGOs.⁶

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. This framework underlines the complementarity between intrinsic motivation and weak accountability in promoting the performance of developmental programs.

In addition to these empirical implications, the model can also provide insights into strengthening accountability in both NGOs and government agencies. In a setting where learning is essential for social service delivery, the model suggests that strong accountability is more suitable for organizations with unmotivated agents. If it is true that agents in NGOs are on average more motivated than government agents, strengthening accountability in government agencies would have a more substantial positive effect on performance than strengthening accountability in NGOs. In fact, if agents in NGOs are highly intrinsically motivated, strengthening accountability might backfire.

Two Agents

Policy learning often involves multiple agents implementing the same policy in their own jurisdictions. Take China, for example, where local officials implement most policy

⁶Social sanctions and enforcement play a decisive role for accountability. Miguel and Gugerty (2005) studies how an inability to impose social sanctions in diverse communities leads to collective action failures in rural western Kenya.

experiments in different localities.⁷ Likewise, American bureaucracies are replete with examples in which various agencies or branches within an agency carry out the same task in their own jurisdictions.⁸ Compared to an agent in a single environment, are agents in these environments more likely to be held accountable for policy outcomes in their own jurisdictions?

To shed some light on this question, I compare the case where two agents implement the same policy in their own jurisdictions with the one-agent case in the main analysis. In the two-agent setting, each agent implements the initial policy in his jurisdiction. As in the one-agent case, the policy outcome in each jurisdiction depends on whether the policy is correct and the agent's effort in that jurisdiction. Unlike the one-agent case, players learn which policy is right based on the outcomes of the initial policy in both jurisdictions. Using this information, the principal makes a unified policy decision for both jurisdictions. Agents implement the policy in their jurisdictions.

Suppose the principal holds an agent accountable for the failure of the initial policy in his own jurisdiction. In that case, policy outcome in the other jurisdiction does not affect whether an agent can stay in the office. Moreover, a successful initial policy in one's jurisdiction reveals that it is correct regardless of its outcome in the other jurisdiction. If the initial policy is successful and the agent stays in office, policy decision and thus the agent's future policy payoff does not depend on the policy outcome in the other jurisdiction. Because policy outcome in the other jurisdiction does not affect the agent's chance of staying in office and agent's policy payoff in future office, an agent does not consider other's initial effort when he decides how much effort to put into the initial policy. In other words, an agent who is held accountable for the initial policy outcome in his own jurisdiction in the two-agent environment puts the same level of

⁷See Cao et al. (1999), Fewsmith (2013), Heilmann (2008), Wang (2009), and Xu (2011) for discussion on policy experimentation on various issues in China.

⁸For example, Bendor (1985) discusses issues in welfare policy in the 1960s.

effort as an agent in the one-agent case.

If an agent is not held accountable for the performance of the initial policy in his own jurisdiction, he stays in office and benefits from policy learning. As shown in the analysis of the main model, an initial effort benefits an agent who faces no accountability through two mechanisms. First, an agent's initial effort reduces the chance of adopting the wrong policy when the initial policy is, in fact, correct. Second, the initial effort improves the learning through failure and helps the agent calibrate his effort better. In the two-agent environment, the beliefs over which policy is the right policy are updated through policy outcomes in both jurisdictions. As long as one jurisdiction observes a success of the initial policy, all players infer that the initial policy is the right policy. As the other works harder on the initial policy, an agent's own effort becomes less crucial for making the right policy decision.⁹ Therefore, compared to an agent in a one-agent setting, an agent in a two-agent environment has incentives to free-ride the other's contribution to the correct decision-making. If the initial policy fails in both jurisdictions, it could be caused by an incorrect policy or by an insufficient effort of both agents. The failure of the initial policy is more informative as an agent works harder on the initial policy. Crucially, as the other agent works harder on the initial policy, the effect of an agent's initial effort on learning through failure decreases.¹⁰ This implies another free-riding incentive of an agent in the two-agent environment. Due to the free-riding incentives, an agent who faces no accountability in a two-agent environment puts lower effort than an agent who is not held accountable for his initial performance in a one-agent setting.

⁹Formally, let e_{11} and e_{21} be the initial effort of the two agents. The ex ante probability of adopting the initial policy when the initial policy is right is $\frac{1}{2}(1 - (1 - e_{11})(1 - e_{21}))$. The effect of an agent's effort e_{11} on this probability is diminishing in the other agent's effort e_{21} .

¹⁰Formally, denote the posterior belief that the alternative policy is right policy upon the failure of the initial policy in both jurisdiction by $\rho(e_{i1}, e_{-i1})$. $\rho(e_{i1}, e_{-i1}) = 1 - \frac{1/2(1-e_{i1})(1-e_{-i1})}{1/2(1-e_{i1})(1-e_{-i1})+1/2}$. $\partial^2 \rho(e_{i1}, e_{-i1}) / \partial e_{i1} \partial e_{-i1} = ((1 - e_{i1})(1 - e_{-i1}) - 1) / ((1 - e_{i1})(1 - e_{-i1}) + 1)^2 < 0$.

Generally speaking, the existence of the other agent does not affect the incentives of an agent who faces accountability but introduces free-riding problems when an agent is not held accountable for his initial performance. The principal who would not have chosen accountability in a one-agent environment now adopts accountability in a two-agent case.

Conclusion

This paper aims to explore a principal's decision whether to hold her agent accountable when learning about the right policy. Accountability links good performance to office-holding and thus motivates agents to exert effort. At the same time, if agents are held accountable for initial failures, agents can't benefit from the learning through failure and therefore hold back initial effort. When agents are highly intrinsically motivated, the principal is less likely to introduce accountability. These ideas are relevant to the discussion of organizations in which agents have an intrinsic preference for performing well on the task. 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 and to understand how the interaction between intrinsic motivations and personnel management affects learning and innovation in the private sector.

References

Acemoglu, D., Kremer, M., and Mian, A. (2008). Incentives in markets, firms, and governments. *Journal of Law, Economics, and Organization*, 24(2):273–306.

- Akerlof, G. A. and Kranton, R. E. (2005). Identity and the economics of organizations. *Journal of Economic perspectives*, pages 9–32.
- Alesina, A. and Tabellini, G. (2007). Bureaucrats or politicians? part i: A single policy task. *The American Economic Review*, pages 169–179.
- Benabou, R. and Tirole, J. (2003). Intrinsic and extrinsic motivation. *The review of economic studies*, 70(3):489–520.
- Bénabou, R. and Tirole, J. (2006). Incentives and prosocial behavior. *American economic review*, 96(5):1652–1678.
- Bendor, J. B. (1985). *Parallel Systems: Redundancy in Government*. University of California Press.
- Bergemann, D. and Hege, U. (2005). The financing of innovation: Learning and stopping. *RAND Journal of Economics*, pages 719–752.
- Besley, T. and Ghatak, M. (2005). Competition and incentives with motivated agents. *American economic review*, 95(3):616–636.
- Besley, T. and Ghatak, M. (2006). Sorting with motivated agents: implications for school competition and teacher incentives. *Journal of the European Economic Association*, 4(2-3):404–414.
- Besley, T. J. and Ghatak, M. (2001). Government versus private ownership of public goods. *The Quarterly Journal of Economics*, 116(4):1343–1372.
- Bolton, P. and Harris, C. (1999). Strategic experimentation. *Econometrica*, 67(2):349–374.
- Bonatti, A. and Hörner, J. (2011). Collaborating. *The American Economic Review*, 101(2):632.
- Brehm, J. O. and Gates, S. (1999). *Working, shirking, and sabotage: Bureaucratic response to a democratic public*. University of Michigan Press.
- Cai, H. and Treisman, D. (2009). Political decentralization and policy experimentation.

- Quarterly Journal of Political Science*, 4(1):35–58.
- Cao, Y., Qian, Y., and Weingast, B. R. (1999). From federalism, chinese style to privatization, chinese style. *Economics of Transition*, 7(1):103–131.
- Cheng, C. and Li, C. (2019). Laboratories of democracy: Policy experimentation under decentralization. *American Economic Journal: Microeconomics*, 11(3):125–54.
- Dixit, A. et al. (2002). Incentives and organizations in the public sector: An interpretative review. *Journal of human resources*, 37(4):696–727.
- Fewsmith, J. (2013). *The Logic and Limits of Political Reform in China*. Cambridge University Press.
- Gailmard, S. and Patty, J. W. (2007). Slackers and zealots: Civil service, policy discretion, and bureaucratic expertise. *American Journal of Political Science*, 51(4):873–889.
- Halac, M., Kartik, N., and Liu, Q. (2016). Optimal contracts for experimentation. *The Review of Economic Studies*, 83(3):1040–1091.
- Heilmann, S. (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, A. V. (2016). Experimentation and persuasion in political organizations. *American Political Science Review*, 110(01):68–84.
- Keller, G., Rady, S., and Cripps, M. (2005). Strategic experimentation with exponential bandits. *Econometrica*, 73(1):39–68.
- Majumdar, S. and Mukand, S. W. (2004). Policy gambles. *The American Economic Review*, 94(4):1207–1222.
- Manso, G. (2011). Motivating innovation. *The Journal of Finance*, 66(5):1823–1860.
- Maskin, E. and Tirole, J. (2004). The politician and the judge: Accountability in government. *American Economic Review*, 94(4):1034–1054.

- Miguel, E. and Gugerty, M. K. (2005). Ethnic diversity, social sanctions, and public goods in kenya. *Journal of Public Economics*, 89(11):2325–2368.
- Prendergast, C. (2008). Intrinsic motivation and incentives. *American Economic Review*, 98(2):201–05.
- Rose-Ackerman, S. (1980). Risk taking and reelection: Does federalism promote innovation? *The Journal of Legal Studies*, pages 593–616.
- Sliwka, D. (2007). Trust as a signal of a social norm and the hidden costs of incentive schemes. *American Economic Review*, 97(3):999–1012.
- Vivalt, E. et al. (2015). Heterogeneous treatment effects in impact evaluation. *American Economic Review*, 105(5):467–70.
- Wang, S. (2009). Adapting by learning the evolution of china’s rural health care financing. *Modern China*, 35(4):370–404.
- Xu, C. (2011). The fundamental institutions of china’s reforms and development. *Journal of Economic Literature*, pages 1076–1151.

Appendix

Proof. Proposition. The proof is conducted in three steps. □

Step 1¹¹ We first characterize $\bar{\lambda}$. Let the derivative of the agent’s objective function given no accountability be U^L . The derivative of the agent’s objective function given no accountability is

$$\frac{\partial U^L}{\partial e_1} = \frac{1}{2} + \frac{1}{2} \frac{\lambda}{2} + \frac{1}{(2 - e_1)^2} \lambda \frac{3}{4} - \frac{e_1}{\lambda}$$

Now observe the following. First, $\frac{\partial U^L}{\partial e_1} > 0$ at $e = 0$. Second, $\frac{\partial U^L}{\partial e_1}$ is convex in

¹¹This step of proof follows Hirsh(2016).

e_1 . Given these two observations, there are at most two solutions to the first order condition given no accountability. The optimum is either the lower solution or the maximum $e_1 = 1$. When there is one solution it is exactly the optimum, and when there are no solutions the optimum is $e_1 = 1$. Thus, whenever $e_1 = 1$ is not the optimum, the problem is well-behaved.

It must be true that if the problem is well behaved for λ' , then holding the other parameters fixed it is also well behaved at $\lambda'' < \lambda'$. Notice that the cross partial of the objective function in (e_1, λ) is $\frac{1}{4} + \frac{1}{(2-e_1)^2} \frac{3}{4} + \frac{e_1}{\lambda^2} > 0$, so the set of optima must decrease when λ decreases. If the problem is well behaved for λ' but not for $\lambda'' < \lambda'$, then $e = 1$ would be optimal for λ'' but not λ' , a contradiction. Thus, for any profile of parameters the set of λ s.t. the problem is well-behaved for all feasible parameters is also an interval $[0, \bar{\lambda})$. When $\frac{\partial U^L}{\partial e}|_{e=1} = \frac{1}{2} + \lambda - \frac{1}{\lambda} < 0$, which holds iff $\lambda < \frac{\sqrt{17}-1}{4}$, the problem is well-behaved.

Step 2. We then solve the threshold value of intrinsic motivation $\hat{\lambda}$ where both types of incentive structures induce the same level of effort by the agent. It must be that

$$\frac{1}{2} + \frac{1}{2} \left(\frac{\hat{\lambda}}{2} + 1 \right) = \frac{e_1}{\hat{\lambda}}$$

and

$$\frac{1}{2} + \frac{1}{2} \frac{\hat{\lambda}}{2} + \frac{1}{(2-e_1)^2} \hat{\lambda} \frac{3}{4} = \frac{e_1}{\hat{\lambda}}.$$

These two equations imply that $\frac{1}{2} = \frac{1}{(2-e_1)^2} \hat{\lambda} \frac{3}{4}$, $e_1 \in [0, 1]$. There exists a unique $e_1 = 2 - \sqrt{3/2\hat{\lambda}}$. Substitute $e_1 = 2 - \sqrt{3/2\hat{\lambda}}$ into the first equation, we have

$$\frac{1}{2} + \frac{1}{2} \left(\frac{\hat{\lambda}}{2} + 1 \right) = \frac{\sqrt{3/2\hat{\lambda}}}{\hat{\lambda}}.$$

$\hat{\lambda} = .777$ solves the above equation.

Step 3. We define the following.

$$MR^H(\lambda) \equiv \frac{1}{2} + \frac{1}{2}\left(\frac{\lambda}{2} + 1\right)$$

$$MR^L(\lambda) \equiv \frac{1}{2} + \frac{1}{2}\frac{\lambda}{2} + \frac{1}{(2 - e_1)^2}\lambda\frac{3}{4}$$

$$MC(\lambda) \equiv \frac{e_1}{\lambda}$$

Let \hat{e} be the equilibrium effort by an agent with intrinsic motivation $\hat{\lambda}$. Notice that for an agent with intrinsic motivation $\hat{\lambda}$, he puts the same level effort under both type of incentive structures. Let $\lambda' \equiv \hat{\lambda} + \delta$ where $\delta \rightarrow 0^+$. Let e^H be the equilibrium level of effort by an agent with intrinsic motivation λ' given high powered incentive. By definition,

$$MR^H(\lambda')|_{e^H} = MC(\lambda')|_{e^H}$$

Notice that $MR^H(\lambda')$ is constant across level of e . This implies the following.

$$MR^H(\lambda')|_{e^H} = MR^H(\lambda') = \frac{1}{2} + \frac{1}{2}\left(\frac{\hat{\lambda} + \delta}{2} + 1\right) = MR^H(\hat{\lambda}) + \frac{1}{4}\delta = MR^H(\hat{\lambda})|_{\hat{e}} + \frac{1}{4}\delta = MC(\lambda')|_{e^H}$$

Given our definition of $\hat{\lambda}$ and \hat{e} , $MR^L(\hat{\lambda})|_{\hat{e}} = MR^H(\hat{\lambda})|_{\hat{e}}$. Substitute $MR^L(\hat{\lambda})|_{\hat{e}} = MR^H(\hat{\lambda})|_{\hat{e}}$ into the above equation, we have

$$MR^L(\hat{\lambda})|_{\hat{e}} + \frac{1}{4}\delta = MC(\lambda')|_{e^H}.$$

Now, let's examine the relation between $MR^L(\hat{\lambda})|_{e^H}$ and $MR^L(\hat{\lambda})|_{\hat{e}}$. Because the equilibrium effort is increasing in λ , it must be that $e^H > \hat{e}$. In addition, $\partial MR^L(\hat{\lambda})/\partial e_1 = \frac{1}{(2-e_1)^3} \lambda \frac{3}{4} > 0$, $MR^L(\hat{\lambda})$ is thus increasing in e_1 . Therefore, $MR^L(\hat{\lambda})|_{e^H} > MR^L(\hat{\lambda})|_{\hat{e}}$.

Notice that

$$MR^L(\hat{\lambda})|_{e^H} = \frac{1}{2} + \frac{1}{2} \frac{\hat{\lambda}}{2} + \frac{1}{(2-e^H)} \hat{\lambda} \frac{3}{4}$$

$$MR^L(\lambda')|_{e^H} = \frac{1}{2} + \frac{1}{2} \frac{\hat{\lambda} + \delta}{2} + \frac{1}{(2-e^H)} (\hat{\lambda} + \delta) \frac{3}{4}$$

It is obvious that $MR^L(\lambda')|_{e^H} > MR^L(\hat{\lambda})|_{e^H} + \frac{1}{4}\delta$. Because $MR^L(\hat{\lambda})|_{e^H} > MR^L(\hat{\lambda})|_{\hat{e}}$,

$$MR^L(\lambda')|_{e^H} > MR^L(\hat{\lambda})|_{\hat{e}} + \frac{1}{4}\delta$$

Recall that $MR^H(\hat{\lambda})|_{\hat{e}} + \frac{1}{4}\delta = MC(\lambda')|_{e^H}$, we have $MR^L(\lambda')|_{e^H} > MC(\lambda')|_{e^H}$. In the equilibrium where the agent faces no accountability, $MR^L(\lambda')|_{e^L} = MC(\lambda')|_{e^L}$. Because the speed in which $MR^L(\lambda')$ is increasing in e_1 is slower than that of $MC(\lambda')$, $e^L > e^H$.

So we have shown that give the intrinsic motivation $\lambda' > \hat{\lambda}$, the equilibrium effort given no accountability is greater than that given accountability ($e^L > e^H$), where $\hat{\lambda}$ is the level of intrinsic motivation given which both incentive structures result in the same level of effort.

Proof. Off Equilibrium Beliefs

A observes her own action. It is reasonable to suppose that A updates belief 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.

Consider the subgame where no accountability is introduced. Suppose A deviates to e_1 . Following a failure and a success, A forms correct beliefs using Bayes' rule as follows.

$$\rho_{1A}^f = \frac{1/2}{1/2(1 - e_1) + 1/2} \geq 1/2$$

$$\rho_{1A}^s = 0 < 1/2$$

If P forms correct beliefs, she chooses $p_2 = b$ following a failure and $p_2 = a$ following a success. A makes a payoff of $\frac{\lambda(\rho_{1A}^f)^2}{2}$ in period 2 if policy a fails and a payoff of $\frac{\lambda}{2}$ 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_{1P}^f , and a success ρ_{1P}^s . To break ties, I assume that P adopts policy b if he is indifferent between two policies in period 2. Given P 's decision rule in period 2, I classify P 's off-equilibrium beliefs into four cases. \square

1. $\rho_{1P}^f \geq 1/2$ and $\rho_{1P}^s < 1/2$. 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. $\rho_{1P}^f \geq 1/2$ and $\rho_{1P}^s > 1/2$. If policy a fails, P makes the same policy decision as she would have if her beliefs are correct. If policy a succeeds, given the off-equilibrium belief, P adopts policy b in period 2. A knows that policy a is correct. But his judgement won't matter because policy a won't be adopted. A puts an effort of and makes a payoff of $\frac{\lambda(\rho_{1A}^f)^2}{2}$. It is less than what he could have made when the principal forms correct beliefs.

3. $\rho_{1P}^f < 1/2$ and $\rho_{1P}^s < 1/2$. If policy a succeeds, P makes the same policy decision as she would have if her beliefs are correct. If policy a fails, given the off-equilibrium belief, P adopts policy a in period 2. A knows that policy a is incorrect, so he doesn't put effort and makes a payoff of 0. A thus makes less than what he could have when the principal forms correct beliefs.
4. $\rho_{1P}^f < 1/2$ and $\rho_{1P}^s > 1/2$. 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 the payoff he could have received when the principal forms correct belief.

I have shown that for any effort deviating from the equilibrium effort 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.

Now consider the subgame where accountability is introduced. Use the same proof strategy as the subgame where there is no accountability. 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 policy a fails, A makes a payoff of 0 regardless of P 's belief. Second, if policy a succeeds, given A 's effort e_1 , the probability that A staying office is $1/2e_1$ regardless of P 's belief. P 's off-equilibrium belief affects A 's payoff only by affecting A 's net policy payoff in period 2, which has been discussed in subgame with no accountability.