Machiavellian Delegation*

Harry Di Pei

Department of Economics, Massachusetts Institute of Technology

March 24, 2015

Abstract

We show that delegating decision rights to a third party can overcome the negative consequences of reputation concerns and improve social welfare. Delegation is valuable by *making public signals noisy*, and therefore, can mitigate the long run players' incentives in reputation building. Our result explains why politicians can shift public blame by delegating unpopular decisions to agents, and establishes a novel role of delegation, which is in sharp contrast to the bulk of literature where the delegated agent has superior information or the ability to acquire information.

Keywords: delegation, bad reputation, blame shifting, belief invariant equilibrium

JEL Codes: D73, D82, D83, P16

*I am indebted to Juuso Toikka for guidance and encouragement. I thank Daron Acemoglu, Philippe Aghion, Kirill Borusyak, Glenn Ellison, Drew Fudenberg, Robert Gibbons, Christian Hellwig, Bengt Holmström, Johannes Hörner, Jonathan Libgober, Benjamin Marx, Alan Kevin Olivi, Heikki Rantakari, Jean Tirole, Shoshana Vasserman, Muhamet Yildiz, seminar participants in Harvard, MIT, University of Leicester and the 2015 AEA Meeting in Boston for helpful comments. Errors are mine.

1 Introduction

Imagine the mayor of a city campaigning for reelection under a plan to reduce crime. Once in power, she has two types of policies in her arsenal: increasing the size of the police force, and tightening the penal code. The efficacy of each policy depends on the situation in the streets. The voters observe the policy that is enacted by the mayor at the time of reelection, but their expertise in evaluating the appropriateness of a policy is limited.

This generates a potential conflict of interest. If the electorate believes that the mayor might be corrupt and wish to increase the police force regardless of its necessity, then even a congruent mayor is forced to tighten the penal code regardless of its effectiveness, in order to convince the public that she is not corrupt and get reelected. In this case, reputation concerns force a congruent mayor to act against public interest. which is the well-known '*pandering*' problem.[1]

We show that delegating decision rights to a third party can overcome the negative consequences of reputation concerns and improve social welfare. We examine an infinite horizon *bad reputation model*, *à la* Ely and Välimäki (2003), with the innovation that the politician (or principal, she) can delegate policy choices to a subordinate official (or agent, he). Both the principal and the agent are patient, and their preferences (or types) are unknown to the public. At the beginning of each period, the public decides whether or not to support the regime. The principal or her agent (depending on whether decision rights are delegated or not) can make a policy choice only if the regime gains public support. The payoff consequences of policies depend on the state of the world. The agent is *indirectly accountable*, since at the end of each period, the principal can either retain him or replace him with another agent. The public then update its belief about the principal's and the agent's types based on the chosen policies as well as the replacement decisions. As a benchmark, when the principal does not have the option to delegate, and is solely responsible for the policies enacted, her incentives to pander are so strong, that in every equilibrium, the public never finds it worthwhile to support her. This result was shown in Ely and Välimäki (2003), and is known as the '*bad reputation effect*'.

We show that both the principal and the agent can be insulated from public pressure under delegation, which results in higher welfare for all players.[2] Intuitively, when the agent cares about

[1]The negative consequences of reputation concerns, i.e. pandering, has been discussed in Morris (2001), Ely and Välimäki (2003), Maskin and Tirole (2004), etc. in contexts of political and organizational decision making.

[2]Ely, Fudenberg and Levine (2008) characterizes conditions on the distribution of commitment types under which reputation is bad. In contrast, all types of all players are rational in our paper, and delegating decision rights is also not considered in their paper.

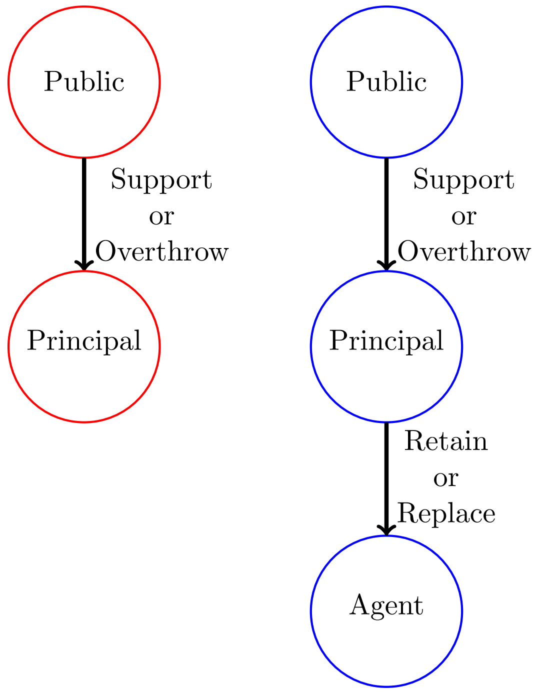

Figure 1: Centralization (left) VS Delegation (Right)

his continuation payoffs, the principal can incentivize him by conditioning her replacement plans on the implemented policy. Since replacement decisions are publicly observable, we also need to ensure that the principal has no incentive to pander (by replacing or retaining an agent who chooses a particular policy). Otherwise, her replacement rule would not be credible. To overcome these difficulties, we construct an equilibrium, in which the policy choices and the replacement decisions are uninformative about the principal's type, and the agent's continuation payoffs are independent of the public's belief about his type. Moreover, the distribution of policies is independent of the agent's type. These conditions together guarantee that the principal is always indifferent between replacing and retaining any agent on the equilibrium path.

The key role of delegation in our model is to make public signals noisy, which is novel relative to the existing literature.[3] In equilibrium, the principal's type is independent of the public signal (i.e. the policy choices and replacement decisions), and it only matters for the correlation between the public signal and the state of the world. The principal and the agent coordinate decisions based

[3]Starting from the seminal contribution of Hölmström (1984), the delegation literature usually assumes that the agent has superior information (for example Dessein [2002], Alonso and Matouschek [2008], etc.), or the ability to acquire information (Aghion and Tirole [1997]). Gibbons, et.al (2012) summarizes the five reasons for delegating decision rights in firms and organizations which has been explored in the literature, and reputation concerns has not been listed.

on their mutual private histories while making the public histories uninformative about the key parameter for the public's payoff—the principal's type.

Our result has two major implications. First, it helps to understand the incentives of indirectly accountable bureaucrats and illustrates the advantage of using them in political decision making. As shown by Maskin and Tirole (2004), both directly accountable politicians and nonaccountable judges have drawbacks: directly accountable officials have pandering incentives, while an official who has judicial power can pursue her private interests without worrying about being replaced. We argue that both the pandering and the congruence problem can be resolved simultaneously by making policy makers indirectly accountable. Under such institutional arrangements, the public signals are noisy. As a result, neither the politician nor the bureaucrat has an incentive to pander since the public learns very little about their intrinsic preferences from the observed policy and personnel choices. The bad bureaucrat's expropriation motives are also limited by the politician's replacement rules.

Second, it helps to explain the wide-spread practice of *blame shifting* in political institutions, where elected officials delegate unpopular decisions to independent agencies when facing public pressure. A well-known example was documented by Niccolò Machiavelli:[4][5]

"...When the duke occupied the Romagna, he found it that the country was full of robbery, quarrels, and every kind of violence; and so, wishing to bring back peace and obedience to authority... He promoted Ramiro d'Orco, to whom he gave the fullest power. This man in a short time restored peace and unity with the greatest success. Afterwards the duke considered that it was not advisable to confer such excessive authority, because he knew that the past severity had caused some hatred against him. So, to clear himself in the minds of the people, and gain them entirely to himself, he desired to show that, if any cruelty had been practised, it had not originated with him, but in the natural sternness of the minister. Under this pretence he took Ramiro, and one morning caused him to be executed and left on the piazza at Cesena... The barbarity of this spectacle caused the people to be at once satisfied..."

[4] The following paragraph is excerpted from '*The Prince*' Chapter VII, '*Concerning New Principalities Which Are Acquired Either By The Arms Of Others Or By Good Fortune*' (Machiavelli [1971]).

[5] Aside from politics, delegation for blame shifting is also observed in the business world, where company executives hire consulting firms to make unpopular but necessary decisions, for example firing employees, lowering wages, etc. For example, company executives hire McKinsey because of the *reputational effect* that hiring McKinsey affords them. Citing a passage from "The Firm: The Story of McKinsey and Its Secret Influence on American Business," a chronicle of McKinsey's business by the journalist Duff McDonald, the article suggests that a component of McKinsey's product is a "cover needed to make an unpopular decision." According to McDonald, "If, as C.E.O., you felt you needed to cut 10 percent of costs, but didn't feel you were getting buy-in from your employees, the hiring of McKinsey generally got the point across quite clearly."

The blame shifting effect of delegation was also found in lab experiments. For example, Hamman et.al. (2010) reports that principals wish to delegate selfish decisions to agents, rather than doing the dirty job themselves; Bartling and Fischbacher (2012) finds that when being treated unfairly, subjects tend to retaliate the direct decision maker, rather than the person who delegates the decision right.

A potential puzzle arises since the agent can usually collude with, or even subject to control by the principal who delegates the decision rights—if the agent's behavior reflects the motivation of the principal, why should blame be attributed to former but not the latter?

Our result implies that blame shifting is possible under delegation, even when the public forms rational expectations, and the principal has interpersonal authority over the agent. When the principal directly chooses the policy herself, it is inevitable for the public to learn about her type from her policy choices. In contrast, thanks to the obfuscation role of delegation, there exist equilibria in which the public can never learn about the principal's type through the chosen policies once decision rights are delegated.[6] As a result, delegation successfully shifts the public blame from the principal to the agent.

The rest of the paper is organized as follows. Section 2 lays out the model, and presents a benchmark result when the principal does not have the option to delegate. Section 3 presents the main result and discusses the tightness of its conditions. Section 4 examines the role of communication and coordination in Machiavellian Delegation. Section 5 concludes and reviews the related literature.

2 The Bad Reputation Model

In this section, we present the baseline bad reputation model, which is introduced by Ely and Välimäki (2003, henceforth, EV), and use their negative result as a benchmark to our analysis on delegation.

2.1 Primitives

We consider an infinite-horizon model, where time is discrete and indexed by $t = 0, 1, 2...$ An infinitely lived '*principal*' (she) interacts with a sequence of short lived '*public*' (it), each living for

[6] In our main result, when the principal and the agent can engage in bilateral communication, the public learns neither the principal nor the agent's type in equilibrium we construct. When the principal and the agent cannot communicate bilaterally, the public can learn about the agent's type (which is irrelevant for the public's as well as the principal's payoff), although they can never learn about the principal's type.

only one period. At the beginning of period t, the current generation public decides whether to support the principal or not. If the principal does not receive public support, all players' payoffs are 0 in that period. If she receives public support, a binary state variable $\theta_t \in \{\theta^0, \theta^1\}$ is realized. We assume that θ_t is i.i.d. across time, and each realization is equally likely. The principal then makes a policy choice, $a_t \in \{0, 1\}$, and the public's payoff is given by:

PUB	$\theta_t = \theta^0$	$\theta_t = \theta^1$
$a_t = 0$	u	$-w$
$a_t = 1$	$-w$	u

where $0 < u < w < 3u$. Hence, the public's payoff is $\frac{u-w}{2}$ when a_t is chosen independently of θ_t, which is strictly negative.

The principal has two possible types, $\omega_p \in \{G, B\}$, which is either good (G) or bad (B). The stage game payoff of the principal (as a function of her type) is shown below:

G	$\theta_t = \theta^0$	$\theta_t = \theta^1$
$a_t = 0$	u	$-w$
$a_t = 1$	$-w$	u

B	$\theta_t = \theta^0$	$\theta_t = \theta^1$
$a_t = 0$	$-w$	$-w$
$a_t = 1$	u	u

To summarize, the good principal shares the same stage game payoffs as the public: both wish to match the policy with the state. The bad principal, however, strictly prefers policy 1 regardless of the state. Each type of principal maximizes her expected discounted average payoff, $V_{\omega_p}(\delta)$, where $\delta \in (0, 1)$ is the discount factor.

We assume that the policy choice a_t is public information. Only a good principal observes θ_t. ω_p is known by the principal but not by the public, who has prior belief:

$$\Pr(\omega_p = B) = \pi_p$$

At the beginning of every period, the newly born public update its belief about ω_p based on the history of policy choices according to Bayes Rule.[7] Importantly, the public can never observe the state.[8]

[7] The public observes a_t only when the principal receives public support in period t, we drop the notation on the public's support decision to simplify notation.

[8] This also implies that the public never directly observes its payoff nor the payoffs of its predecessors'. This assumption is natural when knowing the payoffs from a counterfactual policy requires knowledge of the underlying state. And as noted by James Madison other writers in political science, that voters are much less informed comparing with the bureaucrats and the politicians, and an individual citizen in a large society has strong temptation to free-ride, and is reluctant to spend time studying the intricacies of politics.

2.2 Centralization Benchmark

Our solution concept is *Renegotiation-Proof Nash Equilibrium*, which is a refinement of Nash Equilibrium introduced in EV:

Definition 1 (RPNE). *A Renegotiation-Proof Nash Equilibrium (RPNE) is a Nash Equilibrium in which the principal receives public support in any on-path history at which she is known to be good.*

Renegotiation-Proofness only requires that the good principal always receives public support whenever her type becomes common knowledge, which is plausible in applications where the public (for example, voters, share-holders, small investors) cannot credibly commit to inefficient punishments. EV's negative result is stated as Proposition 1:[9]

Proposition 1. *When $\pi_p > 0$, the principal never receives public support in any RPNE when δ is close enough to* 1.

Proposition 1 implies that all players receive their minmax payoffs, which is 0, in every RPNE when the principal is sufficiently patient. Intuitively, a patient principal cares more about her continuation payoff than her stage game payoff, and the former depends on the public's belief about her type. Since the bad principal always has a stronger incentive to implement policy 1, choosing the alternative policy (policy 0) helps the good principal to distinguish herself from the bad type. When δ is very close to 1, the good principal's reputation concern is so strong that she will invest in it even at the cost of short term public welfare, i.e. she will implement policy 0 even when the state is θ^1. Anticipating this, the public will not support the principal.

As a result, the *bad reputation effect* persists as long as there is a grain of doubt on the principal's type: the good principal cannot use her expertise to serve the public even though her interest is aligned with the public's and latter knows that with arbitrarily high probability.

3 Bad Reputation with Delegation

In this section, we show that having the option to delegate decision rights helps the principal to overcome the bad reputation effect and improve social welfare. We introduce the novel insight that delegation adds noise to public signals, and is valuable even when the agent has no informational advantage over the principal.

[9] This is Theorem 3 in Ely and Välimäki (2003).

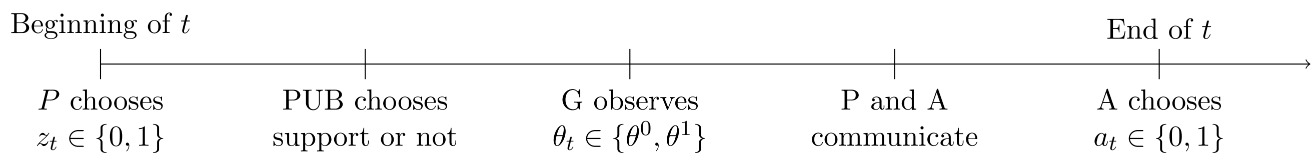

Figure 2: Timeline

3.1 Setup

There is an abundant supply of infinitely-lived agents. Each agent has two possible types, $\omega_a \in \{g, b\}$, which is either good (g) or bad (b), and is identically and independently distributed across agents. Every type of agent receives payoff 0 in every period in which he is not in office. When an agent is in office, his stage game payoff is given by:

g	$\theta_t = \theta^0$	$\theta_t = \theta^1$
$a_t = 0$	u	b
$a_t = 1$	b	u

b	$\theta_t = \theta^0$	$\theta_t = \theta^1$
$a_t = 0$	b	b
$a_t = 1$	u	u

with $b < u$. Hence, the good (or bad) agent shares the same ordinal preference over policies as the good (or bad) principal. Every agent maximizes his expected discounted average payoff, with the same discount factor δ as the principal.

At the beginning of period t, the principal makes her delegation choice $z_t \in \{0, 1\}$: either delegating decision rights to the same agent as in the previous period ($z_t = 1$), or replacing him with a new agent ($z_t = 0$).[10] Once an agent is replaced, he can never resume power again. The public updates its belief about the principal's and agent's types, and decides whether to support the regime or not. If the regime does not receive public support, all players' payoffs are 0 in that period. If the regime receives public support, a binary state variable $\theta_t \in \{\theta^0, \theta^1\}$ is realized and is *only* observed by the good principal. The principal and the agent can exchange cheap talk messages m_t between themselves before the agent making a policy choice $a_t \in \{0, 1\}$. Then the game moves on to the next period. This timeline is summarized in Figure 2:

The public *only* observes the history of policy choices and delegation decisions. Formally, let $x_t \equiv (z_t, a_t)$ be the publicly observable outcome in period t, and let $h^t \equiv (x_0, ..., x_{t-1}, z_t) \in H^t$ be the public history, which is the public's information when it decides whether to grant the regime

[10] For illustration simplicity, we rule out the case where the principal chooses to centralize decision making. Although our results remain valid when centralization is also allowed.

support or not.[11] Let

$$\Pr(\omega_p = B) = \pi_p, \quad \Pr(\omega_a = b) = \pi_a$$

be the public's prior belief about the principal's and the empowered agent's type, and let

$$\pi_p(h^t) \equiv \Pr(\omega_p = B|h^t) \quad \pi_a(h^t) \equiv \Pr(\omega_a = b|h^t)$$

be the public's posterior beliefs at time t, which follow Bayes Rule for every h^t on the equilibrium path.

We assume both the agent and the principal observe ω_p as well as the history of cheap talk messages, $m^t \equiv \{m_0, ..., m_t\}$, in addition to h^t. The principal can only observe the type of an agent *after* she delegates him the decision right. A good principal also observes $(\theta_0, ..., \theta_t)$ in addition to the agent's or a bad principal's private histories. The strategy of a long run player is a function of his (her) private history. The principal's strategy, σ_P, consists of a *replacement plan* and a *communication plan*. The agent's strategy, σ_A, consists of a *policy choice plan* and a *communication plan* in periods when he is delegated the decision right.

3.2 Main Result

Our first result shows that having the option to delegate helps the principal to overcome the bad reputation effect and improve social welfare.

Proposition 2. *When $b \geq -u$ and $\pi_p \leq \frac{3u-w}{2u}$, there exists $\overline{\delta} < 1$, such that for every $\delta \geq \overline{\delta}$, there exists a RPNE in which:*

- *The regime always receives public support;*
- *The good principal's discounted average payoff is at least $\frac{3u-w}{2}$.*

Since $3u > w$, $\frac{3u-w}{2}$ is strictly positive. Hence, according to Proposition 2, the good principal can obtain a strictly positive payoff in equilibrium when the long run players are sufficiently patient and the public is not too pessimistic about the principal's type. This result is not obvious because delegation does not relax the incentive and feasibility constraints in an apparent way. Some problems are listed below:

[11] As in the centralization case, the public cannot observe a_t if the regime does not receive public support in period t. We dispense the notation on the public's decision for simplicity of illustration.

- The agent has no informational advantage over the principal, and hence, when the principal is bad, a_t must still be independent of θ_t.
- The principal can manage to build up her reputation even under delegation. Specifically,
  - Since her replacement decisions are publicly observable, she can potentially affect the public's belief about her type by replacing an agent who implements a particular policy, or an agent who has a bad public reputation.
  - Since information about θ_t is transmitted via cheap talk, the good principal can influence the agent's policy choice via her messages. Hence, if a particular policy results in a better reputation, the good principal has an incentive to mislead the agent at the communication stage. The public will not grant the regime support if it anticipates that the cheap talk messages are uninformative about θ_t, and if this is the case, the bad reputation effect in EV will persist.
- The agent may also have an incentive to pander the public by choosing policy 0 if the principal has an incentive to replace an agent with a bad public reputation.

We explain our idea to overcome these difficulties before presenting our proof. First, since the agent is patient and cares about his continuation payoff, the principal can condition her replacement plans on the implemented policy to incentivize the agent. In order to correct the agent's incentives, we also need to ensure that the principal has the incentive to carry out the equilibrium replacement plan and to communicate θ_t truthfully, which is difficult due to the reasons we have just mentioned. To achieve this, we construct a '*belief invariant equilibrium*', under which the principal can never affect the public's belief about her type:

Definition 2. $(\sigma_P, \sigma_A, \pi_p(h^t), \pi_a(h^t))$ *is a Belief Invariant Equilibrium (BIE) if it is a Sequential Equilibrium, and* $\pi_p(h^t) = \pi_p$ *for every* h^t *on the equilibrium path.*

To achieve belief invariance, we construct strategies under which the unconditional distribution of (z_t, a_t) is independent of ω_p, and ω_p only matters for the correlation between (z_t, a_t) and θ_t. Moreover, the agent's continuation payoff is independent of the public's belief about his type. Aside from being a credible commitment device, another significant role of delegation is to *garble public signals* and to *stop public learning*. This insulates both the principal and the agent from public pressure. As a result, neither of them has any incentive to pander the public.

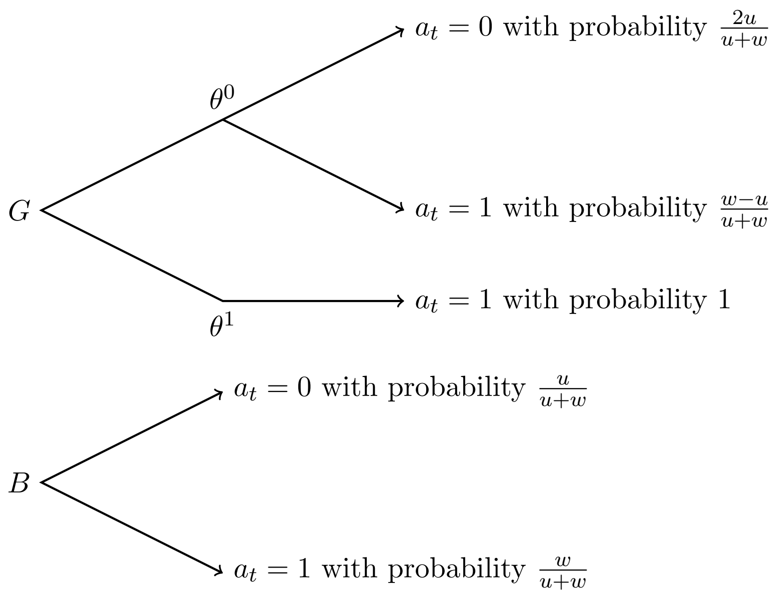

Figure 3: Policy Choices

Furthermore, our result is robust even when the principal cannot directly observe the agent's type. Since equilibrium strategy we display is stationary, it only requires that the agent is sufficiently patient, but does not rely on the principal's discount factor.

Proof of Proposition 2: When $\pi_p = 0$, the claim is trivial. When $\pi_p \neq 0$, belief invariance implies that $\pi_p(h^t) \neq 0$ for every h^t on the equilibrium path. Hence, every BIE satisfies renegotiation proofness, and is also a RPNE.

We describe a BIE which satisfies the requirements in Proposition 2. In our equilibrium, players' strategies only depend on the principal's type, the state of the world, the cheap talk messages and the public's prior belief about the principal's type, but do not depend on the agent's type or the public's belief about the agent's type.

Policy Choice & Communication For every (ω_p, ω_a), policy 0 is chosen with probability $\frac{u}{u+w}$ (unconditional on θ_t) in every period the regime receives public support. The good principal always uses truthful strategies to communicate information about θ_t:

- When $\theta_t = \theta^0$, policy 1 is chosen with probability $\frac{w-u}{u+w}$;
- When $\theta_t = \theta^1$, policy 1 is chosen for sure;

When the principal is bad, the agent's action is independent of θ_t. These probability distributions are displayed in Figure 3.

When $\omega_p = G$ and $\theta = \theta^0$, the principal and the agent use a jointly-controlled lottery (Krishna and Morgan [2004]) in the communication stage to generate a public randomization with two outcomes, with probabilities $\frac{2u}{u+w}$ and $\frac{w-u}{u+w}$.[12] When $\omega_p = B$, the principal and the agent use a jointly controlled lottery to generate a public randomization with two outcomes, with probabilities $\frac{u}{u+w}$ and $\frac{w}{u+w}$. In both cases, the agent is required to choose policy 0 under the first outcome, and policy 1 under the second outcome. If the agent is obedient, then he is retained for sure; if he is disobedient, he is replaced for sure.[13]

Incentive Constraints First, let us check the agent's incentive constraints:

- When $(\omega_p, \omega_a) = (B, g)$, the good agent is indifferent between the two policies, thus has the incentive to choose the required policy.
- When $(\omega_p, \omega_a) = (B, b)$, the bad agent's continuation payoff is $\frac{u(b+w)}{u+w}$, which is greater than his gain from deviating: $\frac{1-\delta}{\delta}(u - b)$, when δ is close enough to 1;
- When $(\omega_p, \omega_a) = (G, g)$, the good agent's continuation payoff is $\frac{3u^2-ub+uw+bw}{2(u+w)}$, which is greater than his gain from deviating: $\frac{1-\delta}{\delta}(u - b)$, when δ is close enough to 1;
- When $(\omega_p, \omega_a) = (G, b)$, the bad agent's continuation payoff is $\frac{u(b+w)}{u+w}$, which is greater than his gain from deviating: $\frac{1-\delta}{\delta}(u - b)$, when δ is close enough to 1;

What remains to be checked is the good principal's incentive constraint at the communication stage. Since the good principal's future payoff is always $\frac{3u-w}{2}$ on the equilibrium path, which is unaffected by her cheap talk messages, and policy 0 is implemented with higher probability when she reports $\theta = \theta^0$ than when she reports $\theta = \theta^1$ and vice versa, she gets a higher payoff in the current period by reporting θ_t truthfully, and thus always has the incentive to do so. □

[12] Krishna and Morgan (2004) show that any rational probabilities can be generated via jointly controlled lotteries when players can engage in bilateral communication. For example, when the required probabilities are $(\frac{1}{3}, \frac{2}{3})$, it can be implemented by letting each player choosing a message in $\{1, 2, 3\}$ simultaneously. The first outcome is implemented if and only if the two players' messages are the same. Each player choosing every message with equal probability is an equilibrium. A notable feature is that no player can unilaterally influence the outcome of the lottery given the other player's mixed strategy.

[13] Notice that when $\omega_p = G$ and $\theta = \theta^1$, the agent is required to choose policy 1 for sure, and no jointly controlled lotteries are needed.

Remark: Proposition 2 shows the existence of equilibrium under delegation such that first, the public histories are uninformative about the principal's type; and second, the agent's continuation value is independent of the public's belief about his type. This institutional arrangement simultaneously eliminates the principal's as well as the agent's pandering incentives, which helps to overcome the *bad reputation effect* and improves social welfare. We call this '*Machiavellian Delegation*'.

Our result sheds light on the long standing puzzle in political science that elected politicians tend to delegate decision rights to independent agencies when they face public pressure and need to make policy choices that are necessary but unpopular. An interesting fact is: delegation can shift public blame even though these agencies are subject to controls by the elected politicians, i.e. they are never really independent.

For example, numerous independent regulatory agencies (IRA) are created in Western European Countries during the 1990s in response to the massive wave of privatization as well as increasing demands from the European Union.[14] These IRAs are responsible for unpopular policy choices such as regulating prices for utilities, regulating competition as well as mergers and acquisitions, etc.[15] Although most of the IRAs are organizationally separated from the ministries, and is neither directly elected nor managed by the elected officials, they are still subject to controls by the elected politicians in more subtle forms. As documented by Thatcher (2002): '*all regulatory agencies face continuing controls by elected officials—nominations, annual budget allocations and requirements to report to legislatures, etc.*' One would worry that due to the lack of '*real independence*' of these IRAs, the pandering incentives of the politician will be passed onto the IRA, and the latter's policy choice would reflect the former's motivation. If this is the case, then delegating decision rights to an IRA will make no difference.

Our equilibrium construction rationalizes the blame shifting role of delegation by emphasizing how delegation makes public signals noisy. Delegation enables the principal and the incumbent agent to coordinate on their mutual private histories, while keeping the public histories uninformative about the key variable to the public's payoff—the principal's type. This ensures that the principal has no incentive to mislead the agent nor to pander the public via replacement decisions. The agent's incentives are corrected by a replacement rule which is self-enforceable in the repeated

[14] Examples include the *Commission for Racial Equality* and the *Food Standards Agency* in the UK, the *Commission nationale de l'informatique et des libertés* in France, as well as various agencies across Europe in regulating the operation of markets and promote competition.

[15] Due to the privatization of public utilities, Western Europe has experienced an increase in utility prices, and the IRA partially absorbs the blame for such tariff increases. Mergers and acquisitions usually involve winners and losers, and part of the complaints from the employees who lost their jobs following successful bids are attributed to the regulators, instead of the government who delegated the decision rights.

game, and in equilibrium, he has no incentive to pander the public nor the principal, and will choose the policy according to the outcome at the communication stage.

Importantly, in order for Machiavellian Delegation to work, the principal cannot offer the agent formal incentive contracts which are unobservable to the public—the agent's incentives can only be provided via '*relational contracts*'. Belief invariance cannot be achieved when the principal can privately commit a replacement rule to the agent (i.e. specifies the replacement rule in a formal contract which is enforceable by a third party) since she can always design it to perfectly align the agent's incentive with hers, and thus, the equilibrium outcome under delegation will be the same as in the case of centralization.[16] Mapping back to our examples, delegating decision rights to an IRA can help a politician to shift public blame when her control over the IRA is relational.[17] Lack of commitment power to the IRA indirectly improves the politician's commitment power to the public, which enables her to implement unpopular policies without negatively affecting her reputation.

3.3 Discussions

In this subsection, we discuss the implications of our result and justify our modeling assumptions. We also illustrate how our model can be adapted to various applications in political economy and organizational design.

Necessity of Belief Threshold Proposition 2 characterizes a sufficient condition on the public's prior belief under which delegation can improve social welfare when the long run players are sufficiently patient. The next Proposition show that this condition is also necessary. Formally, let $\overline{V}_G(\pi_p, \delta)$ be the good principal's supremum discounted average RPNE payoff when π_p is the prior belief on the principal's type and δ is the discount factor. We show that whenever the public assigns probability more than $\frac{3u-w}{2u}$ to the principal being bad, the good principal's expected discounted average payoff is arbitrarily close to 0 in every RPNE when she is sufficiently patient.

Proposition 3. *For every* $\pi_p > \frac{3u-w}{2u}$, $\lim_{\delta \to 1} \overline{V}_G(\pi_p, \delta) = 0$.

[16] A similar result is shown in Katz (1991), that game playing agents cannot change the equilibrium play if there exists a formal contract which perfectly align the principal's and the agent's incentives.

[17] The phenomena that formal incentive contract can crowd out relational incentives is also presence in Baker, Gibbons and Murphy (1994). In their model, formal contracts increases players' incentives to renege in relational contracts by increasing their outside options. In our model, however, formal incentive contract crowds out relational incentives due to the nature of multi-lateral relationships. Given that the principal cannot commit to the public, having more commitment power to the agent can exacerbate the lack of commitment problem between the principal and the public, and in equilibrium, the principal can be better off by having no commitment power to both parties. Thus, in muti-lateral interactions, the value of commitment is non-monotone.

The proof is in the Appendix. Intuitively, in order to achieve a strictly positive payoff when δ is very large, the regime needs to gain public support in the long run and the bad principal's IR constraint needs to be satisfied. The former requires that the good principal's agent matching his policy choice with the state of the world, and the latter requires that the bad principal's agent choosing policy 1 with sufficiently high probability. When the public assigns a high probability that the principal being bad, the public's participation constraint requires that the good principal's agent choosing policy 0 more frequently. This implies that in order to satisfy the public's and the bad principal's IR constraints simultaneously, the observed policy choices under the good and the bad principal must be significantly different. As a result, a_t is very informative about ω_p. Due to the good principal's incentive constraint at the communication stage, i.e. she must prefer to match the policy with the state of the world, the difference in her continuation payoff under policy 0 and policy 1 must be arbitrarily small. Since there exists a belief threshold $\pi^* \equiv \frac{2u}{u+w}$, such that the regime can never get public support when $\pi_p > \pi^{**}$, and this belief is reached with positive probability in finite history, the bad reputation problem in EV arises — the good principal's equilibrium payoff is arbitrarily close to 0 when $\delta \to 1$.

BIE Payoff Frontier The equilibrium we construct in Proposition 2 maximizes the good principal's payoff among all BIE. We focus on the payoff of the good principal since it is an increasing function of the discounted average public welfare. Formally, let

$$V_{PUB}(\delta) \equiv \mathbb{E}\Big[\sum_{t=0}^{\infty}(1-\delta)\delta^t U_{PUB}(t)\Big]$$

be our measure of public welfare, where $U_{PUB}(t)$ is the public's stage game payoff in period t. We have the following relationship between $V_G(\delta)$ and $V_{PUB}(\delta)$ in every BIE in which the principal always receive public support:

$$V_{PUB}(\delta) = (1-\pi_p)V_G(\delta) + \pi_p\frac{u-w}{2}$$

In order to understand the set of payoffs that can be achieved under delegation, the next Corollary characterizes the Pareto payoff frontier of the good and bad principal in BIE for any fixed π_p.

Corollary 3.1. *For any $\pi_p \leq \frac{3u-w}{2u}$, the following payoff vectors can be achieved in BIE when*

δ is close enough to 1*:*

$$\Big(V_G(\delta), V_B(\delta)\Big) = \Big(\frac{u-w}{2} + p_0(u+w), u - p_0(u+w)\Big)$$

where $p_0 \in [\frac{w-u}{2(1-\pi_p)(u+w)}, \frac{u}{u+w}]$.

To obtain these payoffs, it is sufficient to modify our original construction by varying the probability that policy 0 is implemented, which we denote by p_0. Specifically, the good principal still communicates θ_t truthfully. Conditional on the principal being good, policy 0 is implemented with probability $2p_0$ when $\theta_t = \theta^0$ (which requires jointly controlled lottery to decide the policy chosen), and policy 1 is implemented for sure when $\theta_t = \theta^1$. Conditional on the principal being bad, policy 0 is implemented with probability p_0. Increasing p_0 decreases the good principal's equilibrium payoff while increases that of the bad principal's. The restriction on p_0 comes from the bad principal's IR constraint and the public's IR constraint.[18] The former requires that $p_0 \leq \frac{u}{u+w}$; the latter imposes a lower bound on p_0, which is increasing with π_p.

Agent's Information Structure & Communication We assume that the agent does not directly observe θ in order to emphasize the role of delegation in making public signals noisy and to isolate this novel effect from the well-known '*informational effect*'. Different from the majority of delegation models where the agent has valuable information for decision making (for example, Hölmström [1984], Alonso and Matouschek [2008], etc.) or has the ability to acquire information (for example, Aghion and Tirole [1997]), the agent in our model has no informational advantage over the principal.

The principal and the agent can engage in bilateral communication, which also serves as a coordination device in addition to transmitting information about θ_t. Thanks to the jointly controlled lotteries, the agent does not need to be indifferent between the two policies when the principal is good and $\theta_t = \theta^0$, or when the principal is bad. As a result, there is no need for replacement on the equilibrium path, and the agent's continuation value is sufficiently large to discipline his short-run opportunistic behavior as long as δ is close enough to 1.

Although the bilateral communication case we have just studied is prevalent in firms, organizations and political institutions, we will also explore an alternative scenario where the agent is not allowed to send messages to the principal, in order to fully understand the role of coordination

[18] We can show that the good principal's IR constraint does not bind, and is implied by the public's IR constraint.

in Machiavellian Delegation. Lack of coordination requires the agent to mix over two different policies at some information sets, which requires replacement to happen on the equilibrium path whenever the agent strictly prefers one policy over the other in the stage game. As a result, a more demanding condition on the agent's payoff is required in order to guarantee that his continuation payoff is large enough for incentive provision.

Remark on Agent's Payoff In our stage game, the agent only shares the same *ordinal preference* over policies with the principal of the same type, but not the same payoff. Since $b \geq -u > -w$, the agent obtains a strictly higher payoff from an unfavorable policy than the principal.[19] This assumption is to make sure that the good agent's continuation payoff is non-negative when working for the bad principal, which implies that every agent has an incentive to stay in power in all circumstances.

We justify this difference in payoff as the private benefit an agent obtains by staying in office. Consider the following payoff for the good and bad agent:

g	$\theta_t = \theta^0$	$\theta_t = \theta^1$
$a_t = 0$	$u + v$	$-w + v$
$a_t = 1$	$-w + v$	$u + v$

b	$\theta_t = \theta^0$	$\theta_t = \theta^1$
$a_t = 0$	$-w + v$	$-w + v$
$a_t = 1$	$u + v$	$u + v$

where $v > 0$ is the agent's fixed benefit for staying in office, as in line with the political economy literature (Persson and Tabellini [2003]). We can rescale this payoff by normalizing the agent's payoff from a favorable policy to be u, which gives:

g	$\theta_t = \theta^0$	$\theta_t = \theta^1$
$a_t = 0$	u	b
$a_t = 1$	b	u

b	$\theta_t = \theta^0$	$\theta_t = \theta^1$
$a_t = 0$	b	b
$a_t = 1$	u	u

where $\frac{u}{b} = \frac{u+v}{-w+v}$. In the Online Appendix, we provide a micro-foundation to this private benefit by allowing for transfers between the principal and the agent.

Remarks on Modeling Assumptions We discuss several assumptions in our model.

- **Public is short-lived:** Instead of assuming the public is short-lived, our results remain valid when the public is a continuum of small, long lived players, where every individual's decision

[19] Notice that this feature of our model does not undermine the key intuition of our result: the reason for payoff improvement is not a direct consequence of this different payoff function. To see this, even when the principal has the same payoff function as the agent's, her expected discounted payoff in every RPNE is still 0 under centralization.

has negligible impact on the aggregate outcome. In applications where every individual small player's payoff does not depend on the actions of other small players (for example, capital investment), every small long run player plays his myopic best response, behaving as if he is a short run player. In applications where only the aggregate action of small players is payoff relevant (for example, voting), we need an extra requirement that no small player is playing weakly dominated strategies in the dynamic game.[20]

- **Bad principal does not know θ:** The assumption that only the good principal knows the state can be micro-founded by assuming that there is an arbitrary small cost, $\varepsilon > 0$, to acquire information about θ_t. The bad principal has no incentive to acquire that information since her payoff is independent of θ_t and the public never observes θ_t. In reality, the good principal can be interpreted as a benevolent politician who invests in her expertise for policy making; while the bad principal does not care about social welfare, and always prefers a policy which results in higher private benefit.

- **Public does not observe payoff:** This assumption is standard in the repeated game literature, and is relevant in applications when the payoff from an unchosen policy can never be observed, and u and $-w$ are *relative* payoffs to the status quo. The public never observes whether it will be better or worse off under an unchosen policy, and hence, can never infer the true state of the world.

- **Overthrowing versus not supporting the principal:** In our model, the principal gets her minmax payoff for only one period when she loses public support. Our results remain valid when the principal can be overthrown once and for all, which is relevant in revolutions, turnovers, coups, etc.

4 Machiavellian Delegation with Unilateral Communication

In order to emphasize the importance of coordination between the principal and the agent in Machiavellian Delegation, we examine the scenario in which only the principal can communicate to the agent but the agent cannot talk back.[21] Without loss of generality, we focus on equilibria in

[20] An assumption here is that only the aggregate action of the public is observable, but not the individual decisions of small players. This assumption is realistic in scenarios like voting, where players only observe the aggregate distribution of votes.

[21] In the online appendix, we allow the principal to transfer money to the agent. We emphasize the signalling role of monetary transfers, and show the complementarity between transfers and delegation. The draft on the case with

which the principal only sends one message to the agent in every period, which we denote by m_t.

Different from the bilateral communication case, m_t can only be used to transmit information about θ_t, but has no further role in coordinating policy choices and replacement decisions. This requires more demanding conditions on the agent's preferences for delegation to improve social welfare, i.e. the agent's private benefit from staying in office needs to be large enough:

Proposition 4. *When* $b > \frac{(w-u)u}{u+3w}$ *and* $\pi_p \leq \frac{3u-w}{2u}$*, there exists* $\overline{\delta} \in (0,1)$ *such that for every* $\delta > \overline{\delta}$*, there exists a BIE in which:*[22]

- *The regime always receives public support;*
- *The good principal's discounted average payoff is* $\frac{3u-w}{2}$*.*

To prove this result, we construct a BIE such that:

1. All players' strategies are independent of the public's belief about the agent's type;
2. The principal's payoff is independent of the incumbent agent's type, and is always indifferent between replacing and retaining any agent;

We describe the equilibrium below, while leaving the mathematical details to the Appendix.

Constructing BIE in Proposition 4: As in the bilateral communication case, the good principal always reports θ_t truthfully to the agent. In every period, policy 1 is implemented with probability $\frac{w}{u+w}$ (unconditional on θ_t). When $\omega_p = B$, the policy choice is independent of θ_t. When $\omega_p = G$, policy 1 is implemented with probability $\frac{w-u}{u+w}$ when $\theta_t = \theta^0$, and is implemented for sure when $\theta_t = \theta^1$.

To ensure that the above policy rule can be sustained in equilibrium, we construct retention probabilities such that the agent has the incentive to make the required policy choice for every $(\omega_a, \omega_p, \theta_t)$, while simultaneously satisfy the belief invariance condition.

Let $q \equiv \Pr(z_t = 1 | \omega_a = g)$ and $q' \equiv \Pr(z_t = 1 | \omega_a = b)$. The retainment probabilities are summarized below, which are also shown in Figure 4 and 5.

- When $(\omega_a, \omega_p) = (g, G)$ and $\theta = \theta^0$, the agent is retained with probability q when he chooses policy 0, and is retained with probability 1 when he chooses policy 1;

transfers is available upon request.

[22] Note that this lower bound for b is strictly positive but is still strictly smaller than u, meaning that the agent's stage game ordinal preferences are preserved.

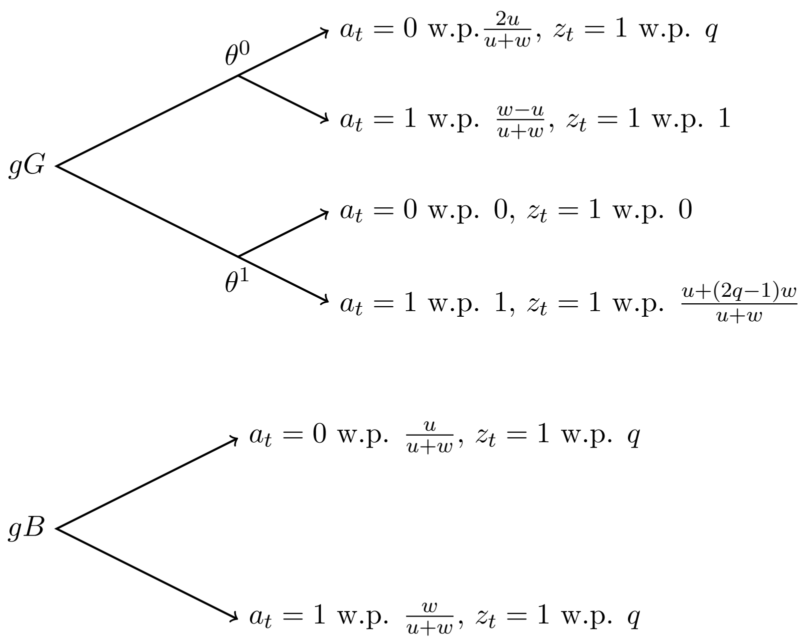

Figure 4: Policy and Replacement Choices When the Agent is Good

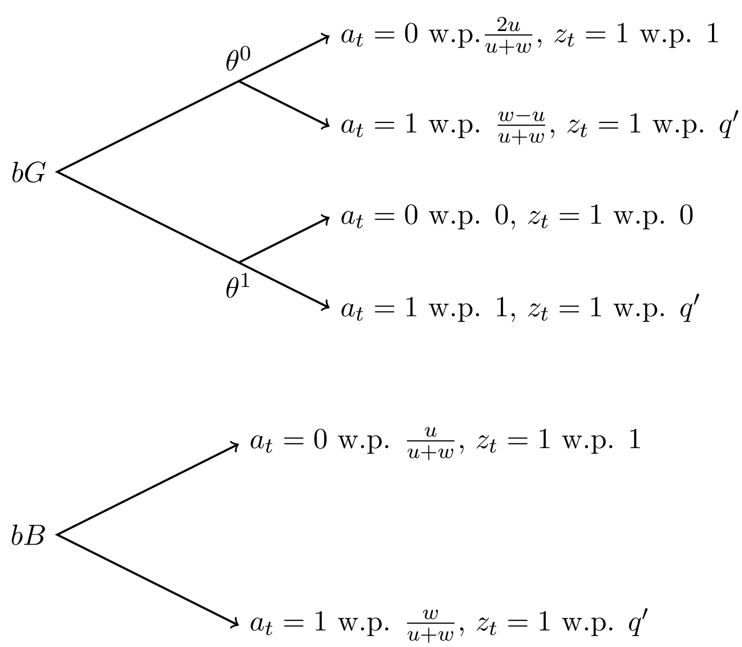

Figure 5: Policy and Replacement Choices When the Agent is Bad

- When $(\omega_a, \omega_p) = (g, G)$ and $\theta = \theta^1$, the agent is retained with probability 0 when he chooses policy 0, and is retained with probability $\frac{u+(2q-1)w}{u+w}$ when he chooses policy 1;
- When $(\omega_a, \omega_p) = (g, B)$, the agent is retained with probability q whichever policy he chooses;
- When $(\omega_a, \omega_p) = (b, G)$ and $\theta = \theta^0$, the agent is retained with probability 1 when he chooses policy 0, and is retained with probability q' when he chooses policy 1;
- When $(\omega_a, \omega_p) = (b, G)$ and $\theta = \theta^1$, the agent is retained with probability 0 when he chooses policy 0, and is retained with probability q' when he chooses policy 1;
- When $(\omega_a, \omega_p) = (b, B)$, the agent is retained with probability 1 when he chooses policy 0, and is retained with probability q' when he chooses policy 1;

It is easy to check that the strategy we described satisfy the belief invariance condition, the principal's incentive constraint in replacement as well as communication. What remains to be checked is the agent's incentive constraint in choosing the required policy. Let $p_0 \equiv \frac{u}{u+w}$ and let V_{ω_a,ω_p} be type ω_a agent's expected discounted payoff when the principal's type is ω_p. V_{ω_a,ω_p} can be expressed as:

$$V_{gG} = \frac{1-\delta}{1-q\delta}\Big[(\frac{1}{2}+p_0)u + (\frac{1}{2}-p_0)b\Big], \quad V_{gB} = \frac{1-\delta}{1-q\delta}\frac{u+b}{2}$$

$$V_{bB} = V_{bG} = \frac{1-\delta}{1-\Big(q'(1-p_0)+p_0\Big)\delta}\Big[(1-p_0)u + p_0 b\Big]$$

where $q\delta$ and $q'\delta$ are the good and the bad agent's *effective discount factor*, respectively.

Every type of agent needs to be indifferent between the two policies whenever he is required to mix, which implies that:

$$(1-\delta)u + \delta q V_{gG} = (1-\delta)b + \delta V_{gG}$$

$$(1-\delta)u + \delta q' V_{bB} = (1-\delta)b + \delta V_{bB}$$

which reduces to:

$$u - b = \delta(1-q')\Big(p_0 b + (1-p_0)u\Big)\frac{1}{1-(q'(1-p_0)+p_0)\delta} \tag{4.1}$$

$$u - b = \delta\frac{1-q}{1-\delta q}\Big(\frac{u+b}{2} + p_0(u-b)\Big) \tag{4.2}$$

What remains to be shown is the following Lemma, the proof of which can be found in the Appendix:

Lemma 4.1. *When $b > \frac{(w-u)u}{u+3w}$, there exists $\overline{\delta} \in (0,1)$ such that for every $\delta > \overline{\delta}$, there exists $q(\delta), q'(\delta) \in (0,1)$ such that (4.1) and (4.2) are satisfied when $q = q(\delta)$ and $q' = q'(\delta)$.*

□

Notice that in the equilibrium we have just constructed, although the public signals are uninformative about the principal's type, but they are informative about the agent's type. The way to insulate both the principal and the agent is to make the agent's type irrelevant for the public's as well as the principal's payoff. Due to the inability to perfectly coordinate policy choices and replacement decisions, replacement happens with positive probability on the equilibrium path. The effective discount factor is bounded away from 1 despite $\delta \to 1$. As a result, the agent's private benefit for staying in office (or his continuation value) needs to be large enough in order for the principal's replacement plan to be effective in deterring the agent's short-run opportunistic behavior.[23]

5 Conclusion & Related Literature

This paper shows that delegating unpopular decisions to an agent can overcome the negative consequences of reputation concerns and improves social welfare. Moreover, it can help the principal to shift public blame and to preserve her reputation in front of the public.[24] We discover a novel role of delegation in obfuscating public learning, which helps to insulate both the principal and the delegated agent from public pressure. In contrast to the informational delegation literature, where the delegated agent has superior information (for example, Hölmström [1984], Dessein [2002], Alonso and Matouschek [2008], etc.) or the ability to acquire information (Aghion and Tirole [1997]),

[23]The belief invariance requirement further exacerbates the '*inefficient replacement*' problem. To see this, $\Pr(z_t = 1|a_t = 0, \theta_t = \theta^0, g, G)$ is strictly below 1 in order to incentivize the good agent to choose policy 1 when the principal is good and $\theta_t = \theta^0$. Belief invariance requires that:

$$\Pr(z_t = 1|a_t = 0, g, B) = \Pr(z_t = 1|a_t = 0, \theta_t = \theta^0, g, G) < 1$$

The incentive of the good agent serving the bad principal further requires that:

$$\Pr(z_t = 1|a_t = 0, g, B) = \Pr(z_t = 1|a_t = 1, g, B) < 1$$

and belief invariance requires that $\Pr(z_t = 1|a_t = 1, \theta_t = \theta^1, g, G) < 1$ in order to ensure that the probability of $z_t = 1$ conditional on $a_t = 1$ is the same when the good agent faces the good and the bad principal.

[24]Aside from EV, the negative consequences of reputation (or career) concerns is also discussed in Hölmström (1999), Morris (2001), Maskin and Tirole (2004), Ottaviani and Sørensen (2006), etc. These papers provide sharp contrasts to the good reputation models, for example, Sobel (1985), that long run incentives causes problem instead of disciplining the players' short run opportunistic behavior.

the agent in our model is neither better informed nor has the expertise to collect information, and the only function of delegation is to garble public signals.

Our result has two major implications. First, it illustrates the advantage for using indirectly accountable bureaucrats in political decision making, and explains how it overcomes the pandering problem associated with directly accountable politicians as well as the congruence problem associated with nonaccountable judges. Second, it rationalizes the '*blame shifting*' role of delegation, and demonstrates its possibility under rational expectations even when the principal has interpersonal authority over the delegated agent.

Related Literature A contemporaneous paper by Vlaicu and Whalley (2014) also uses a hierarchical delegation model to examine the incentives of indirectly accountable policymakers (agents). In their model, the agent observes the state and the principal observes the agent's type. They show that the agent can be insulated from public pressure if and only if the public's belief about the his type is sufficiently optimistic, or the public is sufficiently uncertain about the state. The principal prefers to retain an agent of her own type and in equilibrium can do so between elections. In contrast, we show that the public's belief about the agent's type is irrelevant for incentives, and there is no need for the principal to know the agent's type. In equilibrium, the principal does not necessarily prefer an agent who shares her own preference. The key difference between our paper and theirs is whether the principal observes the state: when she observes the state, as in our model, she can use state contingent replacement plans to correct the agent's pandering incentives as long as the latter is sufficiently patient.

Li (2010) examines a model of strategic communication via intermediaries, and also discovers a blame-sharing role for delegation, i.e. when the sender (government, she) communicates to the receiver (public, it) via an intermediary (media, he), part of the public blame is attributed to the media. In her model, every player's continuation payoff is summarized as a function of the public's belief about his (her) type. However, her model neglects the hidden influence of the government over the media.[25] If such an influence exists, the media's continuation payoff can depend on other factors aside from his public reputation. In our model, all players' continuation payoff functions are endogenous, and can capture other channels through which the agent's payoff is affected, including but not limited to his relationship with the principal.

[25]The phenomena of media capture is prevalent in both developing and developed countries, for example, in Italy, Thailand, Russia, Peru, etc. There is room for collusion between the government and the media when the media owner is vulnerable to political pressure, and can be directly affected by governmental policies. These facts have been formalized and studied in Besley and Prat (2006).

Our work is also related to the literature on the negative effect of transparency in corporate governance and politics. Prat (2005) shows that the principal can be worse off when she has access to more information about the agent, and explains the wide-spread institutional arrangement that the disclosure of information is limited to the public.[26] However, some policy choices are inevitably observed by the public, and Prat (2005) does not offer a solution to the pandering problem in these situations. Our result proposes delegation as an institutional tool to correct the '*wrong kind of transparency*' since it increases the opacity of the decision making process, and therefore, helps to insulate the policy makers from public pressure.

Our work is also related to the literature on hierarchical organizations, pioneered by Tirole (1986), where the public can commit to a '*collusion proof*' mechanism, in order to prevent side-contracting between the principal and the agent.[27] A recent paper by Chassang and Ortner (2014) shows that the public can make collusion harder by creating asymmetric information between the principal and the agent. Since the public has full commitment power in these papers, having the possibility of side-contracting always reduces its welfare. We examine an alternative scenario where the public lacks the commitment power. In our setting, '*relational side-contracting*' between the principal and the agent improves public welfare, and the public benefits from more mutual information between the principal and the agent.

Aghion and Jackson (2014) examines a similar problem, in which a leader is reluctant to make risky decisions when unsuccessful risk taking leads to negative inferences about her competence (type). They use a mechanism design approach, and show that mechanisms with a long review phase followed up by granting tenure approximates first best when players are sufficiently patient.[28] Complementary to theirs, our approach maintains the no commitment assumption on the public's side, and examines the possibility of resolving the bad reputation problem when the leader can delegate decision rights.

[26] The idea that transparency handicaps political decision making and reduces efficiency is in sharp contrast to the conventional wisdom in political science (Persson and Tabellini [2002]), that accountability requires clear responsibility and transparency disciplines political officials and improves voters' welfare.

[27] The terminology used in Tirole (1986) is '*principal-intermediary-agent*', while we adopt the terminology we defined before, i.e. '*public-principal-agent*' in our discussions.

[28] Aghion and Jackson (2014) also examine the case where the public can only commit to term limits, and show that it can also improve the public's payoff comparing with the no commitment benchmark.

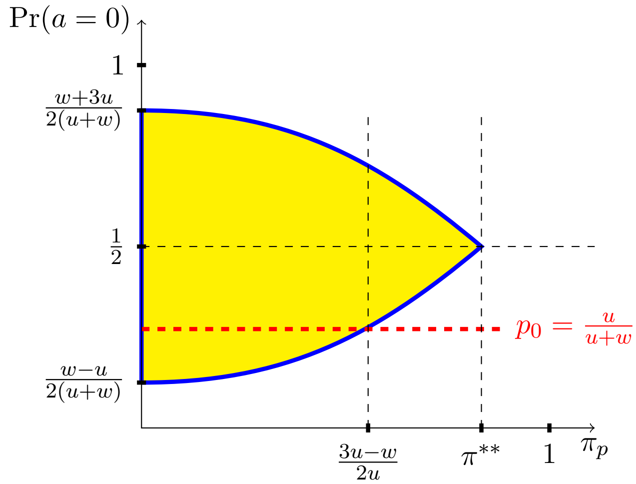

Figure 6: Yellow: Set of Strategies Satisfying Public's IR

A Appendix: Remaining Proofs

Proof of Proposition 3: Let

$$p_0(h^t) \equiv \Pr(a_t = 0|h^t, \omega_p = B), \quad p_1(h^t) \equiv \Pr(a_t = 0|h^t, \omega_p = G)$$

be the probability of implementing policy 0 conditional on the principal's type and the public history.

We list several basic facts. First, the public cannot learn ω_p from a_t at h^t if and only if $p_0(h^t) = p_1(h^t)$. Second, the bad principal's stage game payoff is non-negative if and only if $p_0(h^t) \leq \frac{u}{u+w}$. Third, the public will only support the regime only when the policy choice of the good principal's agent is sufficiently correlated with θ_t,

$$\pi_p(h^t)\frac{u-w}{2} + (1-\pi_p(h^t))\Big(\frac{u-w}{2} + p_1(h^t)(u+w)\Big) \geq 0$$

This implies a lower bound on $p_1(h^t)$:

$$p_1(h^t) \geq \frac{w-u}{2(u+w)}\frac{1}{1-\pi_p(h^t)}$$

Similarly, we can compute an upper bound for $p_1(h^t)$.

$$p_1(h^t) \leq 1 - \frac{w-u}{2(u+w)}\frac{1}{1-\pi_p(h^t)}$$

These two together give the restriction the public's IR imposes on $p_1(h^t)$, which is shown in the yellow region in Figure 6.

Next, we compute two important belief thresholds. First, the public will never support the principal in every history where $\pi_p(h^t)$ exceeds π^{**}, which is defined by:

$$\pi^{**}\frac{u-w}{2} + (1-\pi^{**})u = 0$$

Second, we compute the most pessimistic public belief about ω_p such that there exists $p_0(h^t)$ and $p_1(h^t)$, where the following three conditions are satisfied simultaneously:

- The public's IR;
- The bad principal gets a non-negative payoff in the stage game;
- The public cannot learn ω_p from the policy choice a_t;

Let this belief be π^*, we have:

$$\pi^* \frac{u-w}{2} + (1-\pi^*)\Big((\frac{1}{2} + \frac{u}{u+w})u - (\frac{1}{2} - \frac{u}{u+w}w) \Big) = 0$$

which gives: $\pi^* = \frac{3u-w}{2u}$.

To finish to proof, we use the following learning lemma:

Lemma A.1. *For every* $\{p_0(h^t), p_1(h^t)\}_{h^t}$, $\pi_t(1-\pi_t)||p_{0,t} - p_{1,t}|| \to_p 0$.

The proof can be found in Mailath and Samuelson (2006). Let $(\Omega, \Sigma, \mathcal{F}_t)$ be the probability space over $(\omega_p, \omega_a, x, m)$ under Nash Equilibrium $\Sigma \equiv (\sigma_P, \sigma_A, \pi_p(h^t), \pi_a(h^t))$ and filtration $\mathcal{F}_t$ adapted to the public's information structure. Consider the partition of $\Omega = \Omega_G \cup \Omega_B$, where Ω_G is the event where the principal is good and Ω_B is the event where the principal is bad.

We focus on Ω_B. Since $\pi_p > 0$, so π cannot converge to 0 conditional on $\omega \in \Omega_b$. Hence in the long run, either the principal loses public support for sure (π exceeds π^{**}), or the action distribution is asymptotically uninformative about the principal's type ($||p_0(s^t) - p_1(s^t)|| \to 0$). If the former happens, then according to Ely and Välimäki (2003), the good principal's payoff is close to 0 when $\delta \to 1$. Hence, we only need to focus on the case where $||p_{0,t} - p_{1,t}||$ converges to 0 in probability. Three sets of constraints must be satisfied.

First, the bad principal's individual rationality constraint requires that for any $T \in \mathbb{N}$,

$$\mathbb{E}\Big[\sum_{t=T}^{\infty}(1-\delta)\delta^{t-T}p_0(h^t)\Big] \leq p_0$$

Second, the public's individual rationality constraint needs to be satisfied in every period where the principal receives support. Let $\underline{p}(\pi)$ be the lowest $p_1(h^t)$ which satisfies the public's IR conditional on the public's belief on the principal's type being π, which is shown as the lower bound of the yellow area in Figure 4. A useful observation is that

$$\underline{p}(\pi) = \frac{w-u}{2(u+w)} \frac{1}{1-\pi}$$

which is a strictly increasing, convex function of π.

Third, conditional on $\omega \in \Omega_B$, the public's belief updating process, $\pi(h^t)$, is a sub-martingale, i.e.

$$\mathbb{E}[\pi(h^{t+1})|h^t] \geq \pi(h^t)$$

This together with the convexity of $\underline{p}(\cdot)$ imply that:

$$\mathbb{E}[\underline{p}\Big(\pi(h^{t+1})\Big)|h^t] \geq \underline{p}\Big(\pi(h^t)\Big)$$

which further implies that for every $t_1, t_2 \in \mathbb{N}$ with $t_1 > t_2$,

$$\mathbb{E}\Big[\sum_{t=t_1}^{\infty}(1-\delta)\delta^{t-t_1}\underline{p}\big(\pi(h^t)\big)\Big] \geq \mathbb{E}\Big[\sum_{t=t_2}^{\infty}(1-\delta)\delta^{t-t_2}\underline{p}\big(\pi(h^t)\big)\Big]$$

When $\pi_p > \pi^*$, let

$$\varepsilon = \frac{\underline{p}(\pi_p) - \underline{p}(\pi^*)}{3}$$

If $\lim_{\delta \to 1} \overline{V}_G(\pi_p, \delta) > 0$, Lemma A.1 as well as our argument above implies that there exists $T^* \in \mathbb{N}$ such that for every $T \geq T^*$:

$$\Pr\Big(||p_0(h^T) - p_1(h^T)|| > \varepsilon \Big) < \varepsilon$$

where the probability is only conditional on calendar time. Hence

$$\begin{aligned} p_0 = \underline{p}(\pi^*) &\geq \mathbb{E}\Big[\sum_{t=T}^{\infty}(1-\delta)\delta^{t-T} p_0(h^t)\Big] \\ &\geq \mathbb{E}\Big[\sum_{t=T}^{\infty}(1-\delta)\delta^{t-T} p_1(h^t)\Big] - 2\varepsilon \\ &\geq \mathbb{E}\Big[\sum_{t=T}^{\infty}(1-\delta)\delta^{t-T} \underline{p}\big(\pi(h^t)\big)\Big] - 2\varepsilon \\ &\geq \mathbb{E}\Big[\sum_{t=0}^{\infty}(1-\delta)\delta^{t} \underline{p}\big(\pi(h^t)\big)\Big] - 2\varepsilon \\ &\geq \underline{p}(\pi_p) - 2\varepsilon \end{aligned}$$

which leads to a contradiction. □

Proof of Proposition 4 and Lemma 4.1: We start with defining notation. Let

- q' is the probability of retaining the bad agent when he chooses policy 1;
- q'' is the probability of retaining the bad agent when he chooses policy 0;
- q is the probability of retaining the good agent, which remains unchanged if we only condition on a but not θ;
- q_0 is the probability of retaining the good agent by the good principal when $\theta = \theta^0$ and policy 1 is chosen.

V_{ω_a,ω_p} are defined in the main text. $p_0 \equiv \frac{u}{u+w}$.

We verify that $\{q', q'', q, q_0\}$ can be set to satisfy the agent's incentive constraints for every pair of (ω_p, ω_a):

1. When $(\omega_p, \omega_a) = (B, g)$, the good agent is indifferent between both policies in the stage game, and thus, he has incentive to choose both with strictly positive probability if and only if the probability for continuation is the same across both actions, i.e.

$$q = \Pr(z = 1 | a = 1, g, B) = \Pr(z = 1 | a = 0, g, B)$$

2. When $(\omega_p, \omega_a) = (B, b)$, the bad agent strictly prefers policy 1 to policy 0, and hence, he has incentive to choose both with strictly positive probability if and only if q' and q'' satisfy:

$$(1-\delta)u + \delta q' V_{bB} = (1-\delta)b + \delta q'' V_{bB}$$

3. When $(\omega_p, \omega_a) = (G, b)$, we construct equilibrium such that the constraint in the $(\omega_p, \omega_a) = (B, b)$ case is sufficient. Since policy 0 is only chosen when $\theta_t = \theta^0$, and hence, we can set the probability that the agent is retained under policy 1 to be q' which is the same across both realizations of θ_t, and the probability that the agent is retained under policy 0 is q'' when the state is $\theta_t = \theta^0$ and 0 when the state is θ^1. This gives him indifference when θ^0, and a strictly incentive to choose policy 1 when $\theta_t = \theta^1$.

4. When $(\omega_p, \omega_a) = (G, g)$, the belief invariant condition as well as the stationarity of the equilibrium together imply that:

   $$q = \Pr(z = 1|a = 1, g, G) = \Pr(z = 1|a = 1, g, B) = \Pr(z = 1|a = 0, g, B) = \Pr(z = 1|a = 0, g, G)$$

   The agent needs to be indifferent at θ^0, which requires that:

   $$(1 - \delta)u + \delta q V_{gG} = (1 - \delta)b + \delta q_0 V_{gG}$$

   When $\theta = \theta^1$, the agent is fired for sure if he chooses policy 0 (which happens with 0 probability in equilibrium). Since he strictly prefers policy 1 to policy 0, we only need to make sure that $q_0 < 1$ and

   $$q_0 \leq \frac{q(1 - p)}{\frac{1}{2} - p}$$

   the latter condition guarantees that when $\theta_t = \theta^1$, the probability that the agent is retained is weakly positive. Hence,

   $$q_0 \leq \min\left\{1, \frac{q(1 - p)}{\frac{1}{2} - p}\right\}$$

Plugging in the value function, the simplified constraints are:

$$u - b = \delta(q'' - q')\Big(p_0 b + (1 - p_0)u\Big)\frac{1}{1 - (q'(1 - p_0) + q'' p_0)\delta} \tag{A.1}$$

$$u - b = \delta\frac{q_0 - q}{1 - \delta q}\Big(\frac{u + b}{2} + p_0(u - b)\Big) \tag{A.2}$$

$$q_0 \leq \min\left\{1, \frac{q(1 - p)}{\frac{1}{2} - p}\right\} \tag{A.3}$$

We complete the proof via the following two steps:

- First, we show that for any $b > 0$, there exists $\overline{\delta}$ such that for every $\delta > \overline{\delta}$, there exists $q' \in [0, 1]$ which satisfies (A.1) when $q'' = 1$.
- Second, we show that for any $b > \frac{(w-u)u}{u+3w}$, there exists $\overline{\delta}$, such that for every $\delta > \overline{\delta}$, there exists $q \in (\frac{1}{2}, 1]$, $q_0 > q$, such that (A.2) and (A.3) are satisfied simultaneously.

Step 1: First, when $q'' = q' = 1$,

$$u - b > \delta(q'' - q')\Big(p_0 b + (1 - p_0)u\Big)\frac{1}{1 - (q'(1 - p_0) + q'' p_0)\delta} = 0 \tag{A.4}$$

When $q'' = 1$, $q' = 0$, then,

$$\delta\Big(p_0 b + (1 - p_0)u\Big)\frac{1}{1 - (q'(1 - p_0) + q'' p_0)\delta} > \frac{\delta}{1 - p_0\delta}(1 - p_0)u$$

which is weakly greater than u for δ close enough to 1, and hence, there exists $\overline{\delta} < 1$, such that for every $\delta > \overline{\delta}$,

$$u - b < \delta(q'' - q')\Big(p_0 b + (1-p_0)u\Big)\frac{1}{1-(q'(1-p_0)+q''p_0)\delta} = 0$$

for every $b > 0$. Thus, using Intermediate Value Theorem, for every $b > 0$ and $\delta > \overline{\delta}$, there exists q' such that

$$u - b = \delta(1 - q')\Big(p_0 b + (1-p_0)u\Big)\frac{1}{1-(q'(1-p_0)+q''p_0)\delta} = 0$$

Step 2: First, notice that for every $q > \frac{1}{2}$, $q_0 \leq 1$ is sufficient for (A.3). For any fixed q, when $q_0 = q$,

$$u - b > \delta\frac{q_0 - q}{1-\delta q}\Big(\frac{u+b}{2} + p(u-b)\Big) \tag{A.5}$$

For there to exist q and q_0 which (2) holds, it is sufficient and necessary that:

$$u - b \leq \delta\frac{1 - q}{1-\delta q}\Big(\frac{u+b}{2} + p(u-b)\Big) \tag{A.6}$$

For every $p < \frac{u}{u+w}$,

$$\delta\frac{1 - q}{1-\delta q}\Big(\frac{u+b}{2} + p(u-b)\Big) > \delta\frac{1 - q}{1-\delta q}\Big(\frac{u+\frac{(w-u)u}{u+3w}}{2} + p(u - \frac{(w-u)u}{u+3w})\Big)$$

which converges to $u - \frac{(w-u)u}{u+3w}$ when $\delta \to 1$, which by assumption is greater than $u - b$.

Hence, for every $b < \frac{(w-u)u}{u+3w}$, there exists $\overline{\delta} < 1$ such that every $\delta > \overline{\delta}$, there exists $q_0 \in (q, 1]$ such that (A.2) and (A.3) hold. □

References

[1] Aghion, Philippe, and Matthew O. Jackson (2014) "Inducing Leaders to Take Risky Decisions: Dismissal, Tenure, and Term Limits," Working Paper, Harvard University and Stanford University.

[2] Aghion, Philippe, and Jean Tirole (1997) "Formal and Real Authority in Organizations," *Journal of Political Economy*, 105(1), 1-29.

[3] Alonso, Ricardo, and Niko Matouschek (2008) "Optimal Delegation," *Review of Economic Studies*, 75(1), 259-293. 259-293.

[4] Baker, George, Robert Gibbons and Kevin Murphy (1994) "Subjective Performance Measures in Optimal Incentive Contracts." *Quarterly Journal of Economics*, 109(4), 1125-1156.

[5] Bartling, Björn and Urs Fischbacher (2012) "Shifting the Blame: On Delegation and Responsibility," *Review of Economic Studies*, 79(1), 67-87.

[6] Besley, Timothy, and Andrea Prat (2006) "Handcuffs for the Grabbing Hand? Media Capture and Government Accountability," *American Economic Review*, 96(3), 720-736.

[7] Chassang, Sylvain and Juan Ortner (2014) "Making Collusion Hard: Asymmetric Information as a Counter-Corruption Measure," Working Paper, Princeton University and Boston University.

[8] Dessein, Wouter (2002) "Authority and Communication in Organizations," *Review of Economic Studies*, 69(4), 811-838.

[9] Ely, Jeffrey, Drew Fudenberg and David Levine (2008) "When is Reputation Bad?" *Games and Economic Behavior*, 63, 498-526.

[10] Ely, Jeffrey and Juuso Välimäki (2003) "Bad Reputation," *Quarterly Journal of Economics* 118(3), 785-814.

[11] Gibbons, Robert, Niko Matouschek and John Roberts (2012) "Decisions in Organizations," in Handbook of Organizational Economics, edited by Robert Gibbons and John Roberts, Princeton University Press.

[12] Hamman, John R., George Loewenstein and Roberto A. Weber (2010) "Self-Interest through Delegation: An Additional Rationale for the Principal-Agent Relationship," *American Economic Review*, 100(4), 1826-1846.

[13] Holmström, Bengt, (1984) "On the Theory of Delegation." in M. Boyer and R. Kihlstrom (eds.) *Bayesian Models in Economic Theory* (New York: North-Holland) 115-141.

[14] Holmström, Bengt, (1999) "Managerial Incentive Problems: A Dynamic Perspective," *Review of Economic Studies*, 66(1), 169-182.

[15] Katz, Michael (1991) "Game-Playing Agents: Unobservable Contracts as Precommitments," *RAND Journal of Economics*, 22(3), 307-328.

[16] Krishna, Vijay and John Morgan (2004) "The Art of Conversation: Eliciting Information from Experts through Multi-stage Communication." *Journal of Economic Theory*, 117(2), 147-179.

[17] Li, Wei (2010) "Peddling Influence through Intermediaries," *American Economic Review*, 100(3), 1136-1162.

[18] Machiavelli, Niccolò (1971) "The Prince," Edited and translated by David Wootton, Hackett Publishing Press.

[19] Malaith, George and Larry Sameulson (2006) "Repeated Games and Reputations—Long-Run Relationships," Oxford University Press.

[20] Maskin, Eric and Jean Tirole, "The Politician and the Judge: Accountability in Government." *American Economic Review*, 94(4), 1034-1054.

[21] Morris, Stephen (2001) "Political Correctness," *Journal of Political Economy*, 109(2), 231-265.

[22] Ottaviani, Marco and Peter Norman Sørensen, "Reputational Cheap Talk," *RAND Journal of Economics*, 37(1), 155-175.

[23] Persson, Torsten and Guido Tabellini (2003) "Political Economics: Explaining Economic Policy," MIT Press.

[24] Prat, Andrea (2005) "The Wrong Kind of Transparency," *American Economic Review*, 95(3), 862-877.

[25] Sobel, Joel (1985) "A Theory of Credibility," *Review of Economic Studies* 52(4), 557-573.

[26] Thatcher, Mark (2002) "Delegation to Independent Regulatory Agencies: Pressures, Functions and Contextual Mediation," *West European Politics*, 25(1), 125-147.

[27] Tirole, Jean (1986) "Hierarchies and Bureaucracies: On the Role of Collusion in Organizations," *Journal of Law, Economics and Organizations*, 181-214.

[28] Vlaicu, Razvan and Alexander Whalley (2014) "Hierarchical Accountability in Government: Theory and Evidence," Working Paper, University of Maryland.