

Li, Jin; Matouschek, Niko; Powell, Michael

Working Paper

Power dynamics in organizations

CSIO Working Paper, No. 0139

Provided in Cooperation with:

Department of Economics - Center for the Study of Industrial Organization (CSIO), Northwestern University

Suggested Citation: Li, Jin; Matouschek, Niko; Powell, Michael (2015) : Power dynamics in organizations, CSIO Working Paper, No. 0139, Northwestern University, Center for the Study of Industrial Organization (CSIO), Evanston, IL

This Version is available at:

<https://hdl.handle.net/10419/142013>

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.

POWER DYNAMICS IN ORGANIZATIONS*

Jin Li

Northwestern University

Niko Matouschek

Northwestern University

Michael Powell

Northwestern University

First version: December 2013

This version: December 2015

Abstract

We examine an infinitely repeated game between a principal, who has the formal authority to decide on a project, and a biased agent, who is privately informed about what projects are available. The optimal relational contract speaks to how power is earned, lost, and retained. It shows that entrenched power structures are consistent with optimal administration of power. And it provides new perspectives on why similar firms organize differently, even when those organizational differences lead to persistent differences in performance, and why established firms fail to exploit new opportunities, even when they are publicly observable.

Keywords: internal organization, relational contracting, power

JEL classifications: D23, D82, L23

*We are very grateful for comments from Daniel Barron, Alessandro Bonatti, Wouter Dessein, Matthias Fahn, Daniel Garrett, George Georgiadis, Bob Gibbons, Navin Kartik, Steve Martin, Patrick Rey, Mike Whinston, Yanhui Wu, and participants of various conferences and seminars. We thank Can Urgan for excellent research assistance. All remaining errors are our own. Li: Kellogg School of Management, Northwestern University, jin-li@kellogg.northwestern.edu; Matouschek: Kellogg School of Management, Northwestern University, n-matouschek@kellogg.northwestern.edu; Powell: Kellogg School of Management, Northwestern University, mike-powell@kellogg.northwestern.edu.

1 Introduction

The allocation of formal authority in a firm is carved in stone: it resides with those at the top of the hierarchy and cannot be delegated in a legally binding manner (see, for instance, Bolton and Dewatripont (2013)). Most firms, however, are not autocracies in which a central authority commands all decisions. Instead, those at the top routinely empower their subordinates by promising not to overrule their decisions or to at least exercise restraint in doing so (Baker et al., 1999).

In contrast to formal authority, the allocation of this informal authority or “power” is fluid and changes over time (see, for instance, Pfeffer (1981)). Some of the observed dynamics are intuitive. Divisions, for example, often gain power during periods in which their products are particularly important for the overall profitability of their firms. Other dynamics, however, are more difficult to reconcile with an efficiency-based view of organizations. There are, for instance, many examples in which powerful divisions are able to hold on to their power, even when the peak of their importance has long passed and even when they are using it in ways that are clearly self-serving and harmful to the firm’s profits.

Sears’s catalog division, for example, was able to use the power it had gained during the heyday of mail ordering to delay its closure until the mid-90s, years after analysts started calling for its shut down (Schaefer, 1998). More recently, observers blamed Microsoft’s failure in mobile computing on its powerful Windows and Office divisions. As one news article reported (Eichenwald, 2012):

“Indeed, executives [at Microsoft] said, Microsoft failed repeatedly to jump on emerging technologies because of the company’s fealty to Windows and Office. ‘Windows was the god [...] Ideas about mobile computing with a user experience that was cleaner than with a P.C. were deemed unimportant by a few powerful people in that division, and they managed to kill the effort.’”

In this paper we explore the evolution of power within firms and organizations more broadly. Our goal is to understand whether power can be earned, how it is lost, and why some are able to retain it even when they are using it in openly selfish ways. Existing economic theories of organization are not well-suited to explore these issues, since they are either static or focus on settings in which the optimal allocation of power is stationary (for the former see, for instance, Holmstrom (1984), Aghion and Tirole (1997), and Dessein (2002) and, for the latter, see Baker et al. (1999) and Alonso and Matouschek (2007)). In this paper we therefore expand the existing literature by developing a dynamic model of power.

Our dynamic approach allows us to capture the notion that the prospect of more power tomorrow can motivate subordinates to make good use of whatever power they have today. This notion of power as a form of payment or reward is absent in the economic literature on organizations but

has a long history in sociology and organizational theory. Cyert and March (1963), in particular, observed more than fifty years ago that payments within organizations often take the form of promises about future decisions and decision-making rather than monetary transfers.

As an illustration that will be familiar to many readers, take the dean of a college who wants to convince a department to make a spousal hire. It would be unusual for the dean to try to convince the department to make the hire by promising to pay its members more money. The dean may well, however, promise to reward the department by giving it more discretion in future hiring or to bias future school-wide decisions in its favor. In other words, the dean may pay the department members with power rather than in cash.

The problem with rewarding subordinates by promising them more power in the future is that they value such a promise precisely because it will allow them to bias future decisions in their favor. Paying with power therefore generates a dynamic trade-off between the current and future agency costs of power: it induces subordinates to make good use of whatever power they have today, but it does so only by allowing them to abuse their power in the future.

We show that if those at the top manage this trade-off optimally, then, initially, power is earned and lost in line with an agent's current performance: the agent becomes more powerful if he makes decisions that are good for the organization, and he becomes less powerful if he does not. Eventually, however, this link between power and performance is broken. Depending on the agent's initial performance, the principal then either restricts his power permanently or she permanently expands his power and allows him to make whatever decisions he sees fit. In either case, the organization is no longer able to make efficient use of the agent's expertise, and its performance suffers.

Entrenched power structures, and the frustrations and apparent inefficiencies they entail, therefore need not reflect management failures. Instead, such structures may arise precisely because managers are managing power optimally, albeit in a second-best setting in which their ability to reward employees with money is constrained. These results shed light on dynamic aspects of life inside organizations that have so far received little formal attention. And in doing so, they provide new perspectives on why similar firms organize differently, even though those organizational differences are associated with differences in performance, and why established firms often have a harder time adapting to changes in their environments than their younger rivals, even when the need to adapt to those changes is widely understood. To explore these implications, it is useful to go beyond the broad sketch of intuitions we have given so far and describe our model and results in more detail.

To this end, consider our baseline model, which is an infinitely repeated game between a principal and her agent. Every period, the agent recommends a project, the principal decides which project to choose, and, finally, both parties decide how much effort to put into implementing the project. The principal can choose among multiple projects, including a default project, the agent’s preferred project, and, potentially, the principal’s preferred project. The problem for the principal is that, apart from the default project, only the agent knows which project is which. Moreover, and crucially, the principal’s preferred project is not always available and only the agent knows whether it is. If the agent does not recommend the principal’s preferred project, the principal therefore cannot tell whether the agent is hiding information or the project is simply not available.

Under the optimal relational contract, the principal starts out by promising to rubberstamp the agent’s recommendation. Even though the principal has the formal authority to choose among the projects, the agent, therefore, has the informal authority or “power” to do so. The principal, however, expects the agent to use his power cooperatively. In particular, the principal expects the agent to recommend her preferred project whenever it is available and to recommend the agent’s preferred project only when the principal’s is not available. To motivate the agent to use his power cooperatively, the principal rewards him for recommending her preferred project with an increase in his continuation payoff, and she punishes him for recommending his preferred project with a reduction in his continuation payoff. These changes in the agent’s continuation payoffs capture his future power: the more power the agent is promised in the future, the more often he will be able to choose his preferred project, and the higher his continuation payoff is today.

This initial phase of the relationship continues until the continuation payoff crosses one of two thresholds. If it rises above the upper threshold—because the agent recommended the principal’s preferred project sufficiently often—the principal expands the agent’s power to choose his preferred project. Specifically, the relationship then moves into a steady state in which the principal always rubberstamps the agent’s recommendation, no matter what the agent recommends, and, as a result, the agent always recommends his own preferred project, no matter what projects are available.

If, instead, the continuation payoff drops below the lower threshold—because the agent recommended his preferred project sufficiently often—the principal restricts the agent’s power to choose his preferred project. Specifically, the relationship then moves into one of two steady states: one in which the principal continues to rubberstamp the agent’s recommendation, but only if he does not recommend his own preferred project, and another in which the principal always chooses the default project, no matter what the agent recommends.

In our setting, the relationship therefore does not cycle between punishment and reward phases,

as in Green and Porter (1984). Instead, the principal finds it optimal to delay rewards and punishments for as long as possible before administering them with maximum force. In the long run, the organization can therefore end up with very different allocations of power. Moreover, which allocation the organization ends up with is fully determined by random events during its early history.

These dynamics speak to the debate on whether the organization of firms can be a source of their competitive advantage. Management scholars have long argued that the competitive advantage of some successful firms, such as Lincoln Electric and Toyota, is based on their internal organization. This view has recently been backed up by empirical evidence that differences in managerial and organizational practices are associated with differences in performance (see Bloom et al. (2013) and, for a survey, Gibbons and Hendersen (2013)). Aghion et al. (2015), for instance, find that decentralization is associated with growth in sales, productivity, and profits. In contrast to other sources of competitive advantage, however, a firm’s organization is often well-known, and it is never protected by patents. This raises the question of why, if the competitive advantage of firms is based on their internal organization, under-performing firms don’t simply imitate the organizations of their better-performing rivals.

Our model suggests that the answer to this question may lie in the firms’ pasts. Specifically, the dynamics we described above imply that firms that start out identical can end up with different organizations and thus different performance levels. In particular, some firms will end up as low-performing centralized firms and others will end up as better-performing decentralized ones. In line with the finding in Aghion et al. (2015), therefore, decentralized firms perform better than centralized ones. These organizational and performance differences persist, not because there are informational or legal barriers to imitation. Instead, they persist because firms are constrained by their past promises. The model therefore supports the intuitive view that, as long as there are relational aspects to a firm’s organization, its history can serve as a hidden barrier to imitation.

Another key implication of the dynamics we described above is that the organization gradually gets worse at adapting to changes in the environment. Initially, the principal is able to induce the agent to adapt to those changes in a profit-maximizing manner by making promises about his future power. Eventually, however, the principal has to live up to those promises and either restrict the agent’s power, and thus forgo at least some of the agent’s information, or allow the agent to abuse his power and bias decisions in his favor. In either case, the organization no longer adapts to changes in the environment in a profit-maximizing manner. These dynamics contrast with those in the many models on relationship building in which relationships improve over time, because the

parties learn to cooperate and coordinate with each other (see, for instance, Chassang (2010) and Halac (2014)). In our setting, instead, the relationship deteriorates over time as it gets bogged down by the need to fulfill the very promises that ensured its success early on.

In our baseline model, the changes in the environment that the organization fails to adapt to are privately observed by the agent. The organization’s failure to adapt therefore cannot be detected by an outsider who is simply observing the organization’s current decisions and instead has to be inferred from the equilibrium strategies and the organization’s history. As such, our baseline model does not provide a satisfying explanation for why some firms fail to adapt to new opportunities even when those opportunities are publicly observable, and the failure to adapt is widely critiqued, as in our motivating examples.

To explore the failure to adapt to public opportunities, we build on our baseline model by allowing for a publicly observable opportunity to arrive at some random time during the game. This new project generates a higher expected payoff for the principal than any expected payoff she can realize in its absence. And since its arrival is publicly observable, the principal can simply choose it herself, without having to induce the agent to tell her about it. Nevertheless, we show that if the project arrives late enough, the principal only adopts it with some delay and sometimes does not adopt it at all. The reason for this delay is that the principal now goes beyond promising to tolerate the agent’s biased decision-making by pledging to ignore profitable and publicly observable opportunities. Truly powerful agents therefore do not only get away with making selfish decisions. Their power extends to those at the top of the organization who are bound by their past promises to do the agents’ bidding, even when doing so is known to hurt the organization overall.

These dynamics resonate with the observation that established firms often fail to respond to disruptive changes in their industries. Bower and Christensen (1995, p.43), in particular, observed that “*One of the most consistent patterns in business is the failure of leading companies to stay at the top of their industries when technologies or markets change*” and coined the term “disruptive innovation” to describe this phenomenon. Our model suggests that the very promises that allow firms to adapt to changes in their environments when they are young prevent them from doing so when they are old. The sluggishness of established firms, and the flexibility of their younger rivals, are then two sides of the same coin and are both consistent with optimal management.

2 Literature Review

There is a large literature in sociology and organizational theory on power in organizations. The central question that this literature explores is why some members of organizations wield more

power than others. A common answer is that power is held by those who are deemed to be particularly valuable to the organization. A member may be particularly valuable, for instance, because he controls important resources, as in the “resource dependence theory of power” (Emerson 1962, Pfeffer and Salancik 1978). Or he may be particularly valuable because he helps the organization deal with contingencies, as in the “strategic contingency theory of power” (Hickson, et al., 1971).

The economics literature on power is closely related to, and overlaps with, the incomplete contracting literature on delegation. Most of this literature assumes that the principal can contractually commit to different allocations of decision rights and then explores the formal allocation of those decision rights (see, for instance, Aghion and Tirole (1997), Dessein (2002), and, for a survey, Bolton and Dewatripont (2013)). Courts, however, do not typically enforce contracts between different parties within the same organization (see, for instance, Aghion, et al. (2013) and Bolton and Dewatripont (2013)). In line with this fact, a small number of papers explore the informal allocation of decision rights that arises if those at the top commit to different allocations through non-contractual means. In Aghion and Tirole (1997), for instance, the principal can commit to behaving as if formal authority had been delegated by becoming overloaded and thus staying uninformed. And in Baker, et al. (1999) and Alonso and Matouschek (2007) the principal can commit to behaving as if formal authority had been delegated by agreeing to a relational contract. We follow these last two papers in modeling power as a relational contract. In contrast to those papers, however, we explore a setting in which the optimal allocation of power is not stationary and explore how it changes over time.

Another aspect of our model that we share with those in the incomplete contracting literature on delegation is that we rule out monetary transfers. This assumption captures, albeit in a stark way, the view that the ability of members of an organization to exchange money is often limited by a variety of managerial and legal constraints. The literature on mechanism design without transfers takes this view as its starting point and then explores the optimal design of contracts when parties cannot exchange money but do not face any other constraints on the contracts they can write (see Holmstrom (1984), Melumad and Shibano (1991), and, in a dynamic context, Guo and Horner (2015) and Lipnowski and Ramos (2015))

Our paper is also related to the large literature that studies the economics of relationships; see Samuelson (2006) and Mailath and Samuelson (2006) for reviews. For relationships that survive in the long run, many papers show that their performance improves in general; see Ghosh and Ray (1996), Kranton (1996), Watson (1999, 2002), Mailath and Samuelson (2001), Chassang (2010), Yang (2013), and Halac (2014). Our focus is on the decline of performance among surviving rela-

tionships. One reason for relationships to decline is that the production environment worsens; see, for instance, Garrett and Pavan (2012), and Halac and Prat (2015). The production environment in our model is either stationary—in the baseline model—or improves over time—in the model with public opportunities.

Finally, our paper is related to the literature on dynamic games with one-sided private information; see Mailath and Samuelson (2006) for a general review and Malcomson (1999, 2013) for surveys of applications in labor and organizational economics. In these types of models, the parties use changes in the agent’s continuation payoff to provide incentives. The long-run dynamics then depend on how these continuation payoffs are eventually delivered to the agent. In some models, long-run dynamics involve either termination of the relationship or convergence to an efficient steady state; see, for example, Clementi and Hopenhayn (2006), Biais, et al. (2007), and DeMarzo and Fishman (2007). In other models, punishment is temporary, and relationships forever cycle between punishment and reward phases; see, for example, Padro i Miquel and Yared (2012), Li and Matouschek (2013), Zhu (2013), and Fong and Li (2015). In all these models, the average long-run performance of firms that do not exit is identical. In contrast, in our model, firms can experience long-run organizational and performance differences.

3 The Model

An organization consists of a risk-neutral principal and a risk-neutral agent. Time is discrete and denoted by $t = 1, 2, \dots$. The principal and the agent play the same stage game in every period t . We first describe the stage game and then move on to the repeated game. In the description of the stage game, we omit time subscripts for convenience.

The Stage Game The organization has to decide which project to implement. We follow Mintzberg (1979) and model the decision process as consisting of four stages: the information stage—in which the agent learns which projects are available and what payoffs they generate—the advice stage—in which the agent recommends a project to the principal—the choice stage—in which the principal decides which project to choose—and, finally, the implementation stage—in which the principal and the agent decide how much effort to put into implementing the chosen project. We first describe the projects and then each stage in turn.

Projects: There are up to four projects, which we denote by $n \in \{0, 1, 2, 3\}$. If project n is successful, it pays the agent U_n and the principal Π_n . Project $n = 0$ is the default project and pays $U_0 = \Pi_0 = 0$. Among the remaining three projects, one is the agent’s preferred project, another is the principal’s preferred project, and the final project is a disaster for both parties. For the

agent's preferred project $U_n = \mathbf{B}$ and $\Pi_n = \mathbf{b}$, where $\mathbf{B} > \mathbf{b} > 0$. Analogously, for the principal's preferred project, $U_n = \mathbf{b}$ and $\Pi_n = \mathbf{B}$. Finally, for the disastrous project $U_n = \Pi_n = -\infty$. This payoff structure is a simplified version of the one in Aghion and Tirole (1997). In their model, as in ours, the presence of the disastrous project ensures that both parties prefer the default project to choosing one of the remaining three projects at random. A key feature of our model, and a departure from the payoff structure in Aghion and Tirole (1997), is that the principal's preferred project is only available with probability $p \in (0, 1)$. If the principal's project is "unavailable," it pays $U_n = \Pi_n = -\infty$.

The payoffs U_n and Π_n we just described are the payoffs that project n generates if it is successful. As we will see below, though, projects can also fail and, in particular, they will fail if the parties do not put enough effort into implementing them. If project n does fail, it pays both the agent and the principal the same amount F_n . For the default project, the agent's preferred project, and the principal's preferred project, if it is available, $F_n = U_0 = \Pi_0 = 0$. The assumption that $F_0 = U_0 = \Pi_0$ implies that the default project does not need to be implemented, either because it is a routine project or because it corresponds to "no project." And the assumption that $F_n = U_0 = \Pi_0$ for the other two projects implies that if the parties fail to implement one of those projects, they fall back on the default project. Finally, for the disastrous project and the principal's preferred project if it is not available, we set $F_n = -\infty$. This assumption follows naturally from the fact that these projects pay $-\infty$ even if they are successful.

Information: The payoff from the default project is public information. At the beginning of the stage game, the agent privately observes the payoffs $\theta \equiv \{(U_n, \Pi_n, F_n)\}_{n=1}^3$ of the remaining projects. Each configuration is equally likely. The agent therefore knows the identity of all four projects, while the principal only knows the identity of the default project.

Advice: After the agent learns the payoffs θ , he sends a recommendation $m \in \mathcal{M} \equiv \{0, 1, 2, 3\}$ to the principal. The recommendation is not backed up by any hard evidence and is thus pure cheap talk.

Choice: After the agent has sent his recommendation m , the principal decides which project $k \in \mathcal{K} \equiv \{0, 1, 2, 3\}$ to choose.

Implementation: After the principal has chosen a project, the principal and the agent simultaneously decide how much effort to put into implementing it. The agents's and the principal's efforts are given by $e_A \in \{0, 1\}$ and $e_P \in \{0, 1\}$ and the associated costs are given by ce_A and ce_P , where $c \in (0, b)$. Each party's effort decision is unobserved by the other. The chosen project k is successful if both parties provide effort, and it fails otherwise. As mentioned above, project k

pays the agent and the principal U_k and Π_k if it is successful and F_k if it fails. These payoffs do not include the implementation costs ce_A and ce_P . We will work throughout with the payoffs net of implementation costs, which we denote by $B = \mathbf{B} - c$ and $b = \mathbf{b} - c$.

Randomization: Finally, after the parties have realized their payoffs, they observe the realization $\omega \in [0, 1]$ of a public randomization device, and time moves on to the next period. We summarize the timing of the game in Figure 1.




Figure 1: Timing of the stage game.

The Repeated Game The agent and the principal have a common discount factor $\delta \in (0, 1)$. At the beginning of any period t , the principal's expected payoff is therefore given by

$$\pi_t = (1 - \delta) E_t \left[\sum_{\tau=t}^{\infty} \delta^{\tau-t} (e_{P,t} e_{A,t} \Pi_{k_t} + (1 - e_{P,t} e_{A,t}) F_{k_t} - c \cdot e_{P,t}) \right],$$

and the agent's expected payoff is given by

$$u_t = (1 - \delta) E_t \left[\sum_{\tau=t}^{\infty} \delta^{\tau-t} (e_{P,t} e_{A,t} U_{k_t} + (1 - e_{P,t} e_{A,t}) F_{k_t} - c \cdot e_{A,t}) \right].$$

Note that we multiply the right-hand side of each expression by $(1 - \delta)$ to express payoffs as per-period averages.

We follow the literature on repeated games with imperfect public monitoring and define a relational contract as a pure-strategy Perfect Public Equilibrium (henceforth PPE) in which the principal and the agent play public strategies and, following every history, the strategies are a Nash Equilibrium of the continuation game. Public strategies are strategies in which the players condition their actions only on publicly available information. In particular, the agent's strategy does not depend on her past private information. Our restriction to pure strategies is without loss of generality, because our game has only one-sided private information and is therefore a game with a product monitoring structure. In this case, there is no need to consider private strategies since every sequential equilibrium outcome is also a PPE outcome (see, for instance, p.330 in Mailath and Samuelson (2006)).

Formally, let $h_t = \{m_\tau, k_\tau, e_{A,\tau}, e_{P,\tau}, \omega_\tau\}_{\tau=1}^{t-1}$ denote the public history at the beginning of any period t , and let H_t denote the set of period- t public histories. Note that $H_1 = \emptyset$. A public strategy for the principal is a sequence of functions $\{K_t, E_{P,t}\}_{t=1}^\infty$, where $K_t : H_t \times \mathcal{M}_t \rightarrow \mathcal{K}_t$, and $E_{P,t} : H_t \times \mathcal{M}_t \times \mathcal{K}_t \rightarrow \{0, 1\}$. Similarly, a public strategy for the agent is a sequence of functions $\{M_t, E_{A,t}\}_{t=1}^\infty$, where $M_t : H_t \times \Theta_t \rightarrow \mathcal{M}_t$ and $E_{A,t} : H_t \times \Theta_t \times \mathcal{K}_t \rightarrow \{0, 1\}$.

We define an “optimal relational contract” as a PPE that maximizes the principal’s first-period equilibrium payoff. Our goal is to characterize the set of optimal relational contracts.

4 Preliminaries

In this section, we follow the approach pioneered by Abreu, et al. (1990) to characterize the PPE payoff set. Each equilibrium payoff pair (u, π) can be supported either by randomization among several equilibrium payoff pairs or by a pure action of the stage game and a pair of continuation payoffs associated with each public outcome. Public randomization occurs at the end of the period; at the beginning of each period, the players therefore select a pure action of the stage game, receive the flow payoffs generated by that pure action, and expect to receive particular continuation payoffs. We begin our analysis in Section 4.1 by describing the constraints that have to be satisfied for an equilibrium payoff pair to be supported by a particular action and a particular set of continuation payoffs. We then characterize the PPE payoff set as the solution to a functional equation, which we describe in Section 4.2. Given our explicit characterization of the PPE payoff set, we then describe the dynamics of an optimal relational contract in Section 5.

4.1 The Constraints

We denote the PPE payoff set by \mathcal{E} . Any payoff pair $(u, \pi) \in \mathcal{E}$ is either generated by pure actions or by randomization among equilibrium payoff pairs that are each generated by pure actions. Each set of pure actions corresponds to a different organizational arrangement. We focus on four arrangements. The focus on these four arrangements is without loss of generality since, as we will show below, the optimal relational contract can be sustained without making use of any other arrangement.

Under the first arrangement, the principal chooses the default project, no matter what the agent recommends. Moreover, neither party puts effort into implementing the project since the default project does not require implementation. Since the principal makes no use of the agent’s information under this arrangement, we refer to it as centralization and denote it by C .

Under the other three arrangements, the principal rubberstamps whichever project n the agent

recommends. Even though the principal still has the formal authority to choose a project, she therefore now gives the agent the “informal authority” or “power” to make this choice. The three arrangements, however, differ in the recommendation rule that the principal expects the agent to follow. Under cooperative empowerment E_C the principal expects the agent to recommend the principal’s preferred project if it is available and the agent’s preferred project otherwise. Under restricted empowerment E_R , the principal again expects the agent to recommend the principal’s preferred project if it is available. If it is not available, however, she now expects the agent to recommend the default project rather than the agent’s preferred project. Finally, under unrestricted empowerment E_U , the principal expects, and tolerates, that the agent always recommends his preferred project. Under all three arrangements, the agent behaves as the principal expects him to behave, and both parties put effort into implementing whichever project the principal chooses.

In the remainder of this section, we first discuss the constraints that have to be satisfied for a payoff pair $(u, \pi) \in \mathcal{E}$ to be generated by one of these organizational arrangements. We then conclude the section by stating the constraint that has to be satisfied if the payoff pair is generated by randomization.

Centralization C Under centralization, the principal chooses the default project, no matter what the agent recommends. Given this decision rule, we can assume without loss of generality that the agent simply recommends the default project. A payoff pair (u, π) can be supported by centralization if the following constraints are satisfied.

(i.) Feasibility: For the continuation payoffs to be feasible, they also need to be PPE payoffs. The continuation payoffs u_C and π_C the parties realize under centralization therefore have to satisfy the self-enforcement constraint

$$(u_C, \pi_C) \in \mathcal{E}. \quad (\text{SE}_C)$$

(ii.) No Deviation: To ensure that neither party deviates, we need to consider both off- and on-schedule deviations.

Off-schedule deviations are deviations that both parties can observe. We assume that if an off-schedule deviation occurs, the parties never again implement any projects, and the principal never again chooses a project other than the default project. This assumption is without loss of generality since it is the worst possible equilibrium and gives each party its minmax payoff. Given this punishment rule, neither player has an incentive to deviate off-schedule, since payoffs on the equilibrium path are weakly positive while punishment payoffs are weakly negative.

In contrast to off-schedule deviations, on-schedule deviations are privately observed. There are

no on-schedule deviations under centralization since the agent's recommendation does not depend on his private information and the principal does not have any private information.

(iii.) Promise Keeping: Finally, the consistency of the PPE payoff decomposition requires that the parties' payoffs are equal to the weighted sum of current and future payoffs. The promise-keeping constraints

$$\pi = \delta \pi_C \tag{PK_C^P}$$

and

$$u = \delta u_C \tag{PK_C^A}$$

ensure that this is the case.

Cooperative Empowerment E_C Under cooperative empowerment, the principal rubberstamps the agent's recommendation and the agent recommends the principal's preferred project when it is available and the agent's preferred project otherwise. A payoff pair (u, π) can be supported by cooperative empowerment if the following constraints are satisfied.

(i.) Feasibility: For the continuation payoffs to be feasible, they also need to be PPE payoffs. Let $(u_{E_C, \ell}, \pi_{E_C, \ell})$ denote the parties' continuation payoffs if the agent recommends his preferred project and let $(u_{E_C, h}, \pi_{E_C, h})$ denote their payoffs if the agent recommends the principal's preferred project. The self-enforcement constraint is then given by

$$(u_{E_C, \ell}, \pi_{E_C, \ell}), (u_{E_C, h}, \pi_{E_C, h}) \in \mathcal{E}. \tag{SE_{E_C}}$$

(ii.) No Deviation: The principal and the agent never want to deviate off schedule. Moreover, the principal has no on-schedule deviations. The agent, however, can deviate on schedule by recommending his preferred project when the principal's preferred project is available. The incentive constraint

$$(1 - \delta) b + \delta u_{E_C, h} \geq (1 - \delta) B + \delta u_{E_C, \ell} \tag{IC_{E_C}}$$

ensures that he does not want to do so.

(iii.) Promise Keeping: The promise-keeping constraints are now given by

$$\pi = p [(1 - \delta) B + \delta \pi_{E_C, h}] + (1 - p) [(1 - \delta) b + \delta \pi_{E_C, \ell}] \tag{PK_{E_C}^P}$$

and

$$u = p [(1 - \delta) b + \delta u_{E_C, h}] + (1 - p) [(1 - \delta) B + \delta u_{E_C, \ell}]. \tag{PK_{E_C}^A}$$

Restricted Empowerment E_R Under restricted empowerment, the principal rubberstamps the agent's recommendation and the agent recommends the principal's preferred project when it is available and the default project otherwise. A payoff pair (u, π) can be supported by restricted empowerment if the following constraints are satisfied.

(i.) Feasibility: For the continuation payoffs to be feasible, they also need to be PPE payoffs. Let $(u_{E_R, \ell}, \pi_{E_R, \ell})$ denote the parties' continuation payoffs if the agent recommends the default project and let $(u_{E_R, h}, \pi_{E_R, h})$ denote their payoffs if he recommends the principal's preferred project. The self-enforcement constraint is then given by

$$(u_{E_R, \ell}, \pi_{E_R, \ell}), (u_{E_R, h}, \pi_{E_R, h}) \in \mathcal{E}. \quad (\text{SE}_{E_R})$$

(ii.) No Deviation: The principal never wants to deviate off schedule. The agent can deviate off schedule by recommending his own project. If he does so, he receives $(1 - \delta)B$ this period followed by 0. To prevent the agent from deviating off-schedule, we need that

$$u \geq (1 - \delta)B. \quad (\text{IC}_{E_R}^{\text{Off}})$$

Moreover, the principal has no on-schedule deviations. The agent, however, can deviate on schedule by recommending the default project when the principal's preferred project is available. The incentive constraint

$$(1 - \delta)b + \delta u_{E_R, h} \geq \delta u_{E_R, \ell} \quad (\text{IC}_{E_R}^{\text{On}})$$

ensures that he does not want to do so.

(iii.) Promise Keeping: The promise-keeping constraints are now given by

$$\pi = p[(1 - \delta)B + \delta \pi_{E_R, h}] + (1 - p)\delta \pi_{E_R, \ell} \quad (\text{PK}_{E_R}^{\text{P}})$$

and

$$u = p[(1 - \delta)b + \delta u_{E_R, h}] + (1 - p)\delta u_{E_R, \ell}. \quad (\text{PK}_{E_R}^{\text{A}})$$

Unrestricted Empowerment E_U Under unrestricted empowerment, the principal rubberstamps the agent's recommendation even though the agent always recommends his preferred project. A payoff pair (u, π) can be supported by unrestricted empowerment if the following constraints are satisfied.

(i.) Feasibility: We denote continuation payoffs under unrestricted empowerment by (u_{E_U}, π_{E_U}) . The self-enforcement constraint is then given by

$$(u_{E_U}, \pi_{E_U}) \in \mathcal{E}. \quad (\text{SE}_{E_U})$$

(ii.) No Deviation: As in the case of centralization, the principal and the agent never want to deviate off schedule and there are no feasible on-schedule deviations.

(iii.) Promise Keeping: The promise-keeping constraints are now given by

$$\pi = (1 - \delta) b + \delta \pi_{EU} \quad (\text{PK}_{EU}^P)$$

for the principal and

$$u = (1 - \delta) B + \delta u_{EU} \quad (\text{PK}_{EU}^A)$$

for the agent.

Randomization Finally, a payoff pair (u, π) can be supported by randomization. In this case, there exist at most three distinct PPE payoffs $(u_i, \pi_i) \in \mathcal{E}$, $i = 1, 2, 3$ such that

$$(u, \pi) = \alpha_1 (u_1, \pi_1) + \alpha_2 (u_2, \pi_2) + \alpha_3 (u_3, \pi_3)$$

for some $\alpha_1, \alpha_2, \alpha_3 \geq 0$ and $\alpha_1 + \alpha_2 + \alpha_3 = 1$.

Assumptions on Parameters We now make three assumptions on the parameters of the model.

ASSUMPTION 1. $B/b \leq p/(1 - \delta)$.

ASSUMPTION 2. $B/b \leq 1/p$.

ASSUMPTION 3. If $\delta < 1/2$, $B/b \leq [p(1 - \delta) + \delta(1 - p)] / [(1 - 2\delta)(1 - p)]$.

The first assumption guarantees that both restricted empowerment and cooperative empowerment can be sustained in equilibrium. In particular, it implies that there exist feasible pairs of continuation payoffs $u_{EC,\ell}$ and $u_{EC,h}$ that are sufficiently spread apart so that the agent's incentive constraint is satisfied. The second assumption ensures that unrestricted empowerment is better for the principal than restricted empowerment, and the second and third assumptions ensures that the optimal relational contract involves nontrivial dynamics. As long as $p + \delta \geq 1$, these assumptions can hold simultaneously.

4.2 The Constrained Maximization Problem

We now use the techniques developed by Abreu, et al. (1990) to characterize the PPE payoff set and, in particular, its frontier. For this purpose, we define the payoff frontier as

$$\pi(u) \equiv \sup \{ \pi' : (u, \pi') \in \mathcal{E} \},$$

where \mathcal{E} is the PPE payoff set.

We can now state our first lemma, which establishes several properties of the PPE payoff set. The proofs of this lemma and all other results are in the appendices.

LEMMA 1. *The PPE payoff set \mathcal{E} has the following properties: (i.) it is compact; (ii.) $\pi(u)$ is concave; (iii.) $\inf \{u : (u, \pi) \in \mathcal{E}\} = 0$ and $\sup \{u : (u, \pi) \in \mathcal{E}\} = B$.*

The first part of the lemma shows that the PPE payoff set is compact. This result immediately follows from the assumption that there are only a finite number of actions. It implies that for any $u \in [0, B]$ the payoff pair $(u, \pi(u))$ is in the PPE payoff set. The second part of the lemma shows that the payoff frontier is concave, which follows directly from the availability of a public randomization device. Finally, the third part shows that the agent's smallest and largest PPE payoffs are 0 and B .

We now proceed to characterize the PPE payoff frontier $\pi(u)$. To characterize the frontier, we need to determine, for each $(u, \pi(u)) \in \mathcal{E}$, whether it is supported by a pure action $j \in \{C, E_R, E_C, E_U\}$ or by randomization. Moreover, if it is supported by a pure action j , we need to specify the associated continuation payoffs. The next lemma characterizes the principal's continuation payoff for any of the agent's continuation payoffs, regardless of the actions that the parties take.

LEMMA 2. *For any $(u, \pi(u))$, the continuation payoffs are also on the frontier.*

The lemma shows that payoffs on the frontier are sequentially optimal. This is the case since the principal's actions are publicly observable. It is therefore not necessary to punish her by moving below the PPE frontier. This feature of our model is similar to Spear and Srivastava (1987) and the first part of Levin (2003), in which the principal's actions are also publicly observable. In contrast, joint punishments are necessary when multiple parties have private information as, for instance, in Green and Porter (1984), Athey and Bagwell (2001), and the second part of Levin (2003).

Having characterized the principal's continuation payoff for any of the agent's continuation payoffs in the previous lemma, we now state the agent's continuation payoffs associated with each action in the next lemma.

LEMMA 3. *For any payoff pair $(u, \pi(u))$ on the frontier, the agent's continuation payoffs satisfy the following conditions:*

(i.) *If the payoff pair is supported by centralization, the agent's continuation payoff is*

$$u_C(u) = u/\delta.$$

(ii.) If the payoff pair is supported by restricted empowerment, there exists a payoff-equivalent equilibrium in which the agent's continuation payoffs are

$$u_{E_R,h}(u) = u_{E_R,\ell}(u) \equiv u_{E_R}(u) = (u - (1 - \delta)pb) / \delta.$$

(iii.) If the payoff pair is supported by cooperative empowerment, there exists a payoff-equivalent equilibrium in which the agent's continuation payoffs are

$$u_{E_C,h}(u) = (u - (1 - \delta)b) / \delta$$

and

$$u_{E_C,\ell}(u) = (u - (1 - \delta)B) / \delta.$$

(iv.) If the payoff pair is supported by unrestricted empowerment, the agent's continuation payoff is

$$u_{E_U}(u) = (u - (1 - \delta)B) / \delta.$$

In the cases of centralization and unrestricted empowerment, the agent's continuation payoffs follow directly from the promise-keeping constraints PK_C^A and $PK_{E_U}^A$. In the case of restricted empowerment, the agent's continuation payoffs follow from the promise-keeping constraint $PK_{E_R}^A$ and from setting $u_{E_R,h}(u) = u_{E_R,\ell}(u)$. Setting $u_{E_R,h}(u) = u_{E_R,\ell}(u)$ is optimal, because it satisfies the agent's incentive constraint and, since π is concave, it makes the principal better off. In the case of cooperative empowerment, the agent's continuation payoffs follow directly from combining the promise-keeping constraints with the agent's incentive constraint IC_{E_C} , which we can, without loss of generality, take to be binding. To see that we can do so, suppose the incentive constraint is not binding. We can then reduce $u_{E_C,h}$ and increase $u_{E_C,\ell}$ in such a way that u remains the same, and all the relevant constraints continue to be satisfied. Again, since π is concave, such a change makes the principal weakly better off.

Next we use Lemmas 2 and 3 to provide expressions for the principal's expected payoff for a given action and a given expected payoff for the agent. For this purpose, let $\pi_j(u)$ for $j \in \{C, E_R, E_C, E_U\}$ be the highest equilibrium payoff to the principal given action j and agent's payoff u . We then have

$$\pi_C(u) = \delta\pi(u_C(u)),$$

$$\pi_{E_R}(u) = p(1 - \delta)B + \delta\pi(u_{E_R}(u)),$$

$$\pi_{E_C}(u) = p[(1 - \delta)B + \delta\pi(u_{E_C,h}(u))] + (1 - p)[(1 - \delta)b + \delta\pi(u_{E_C,\ell}(u))],$$

and

$$\pi_{E_U}(u) = (1 - \delta)b + \delta\pi(u_{E_U}(u)).$$

We can now state the next lemma, which describes the constrained maximization problem that characterizes the payoff frontier.

LEMMA 4: *The PPE frontier $\pi(u)$ is the unique function that solves the following problem. For all $u \in [0, B]$*

$$\pi(u) = \max_{\alpha_j \geq 0, u_j \in [0, B]} \sum_{j \in \{C, E_R, E_C, E_U\}} \alpha_j \pi_j(u_j)$$

such that

$$\sum_{j \in \{C, E_R, E_C, E_U\}} \alpha_j = 1$$

and

$$\sum_{j \in \{C, E_R, E_C, E_U\}} \alpha_j u_j = u.$$

The lemma shows that any payoff pair on the frontier is generated either by a pure action j —in which case the weight α_j is equal to one—or by randomization—in which case α_j is less than one. We obtain the frontier by choosing the weights optimally.

5 The Optimal Relational Contract

In this section we characterize the optimal relational contract, that is, the PPE that maximizes the principal's expected payoff. For this purpose, we first characterize the payoff frontier by solving the constrained maximization problem in Lemma 4.

LEMMA 5. *There exist two cut-off levels $\underline{u}_{E_C} \in [(1 - \delta)B, (1 - \delta)B + \delta pb]$ and $\bar{u}_{E_C} = (1 - \delta)b + \delta B$ such that the PPE payoff frontier $\pi(u)$ is divided into four regions:*

(i.) *For $u \in [0, pb]$, $\pi(u) = Bu/b$ and $(u, \pi(u))$ is supported by randomization between centralization and restricted empowerment.*

(ii.) *For $u \in [pb, \underline{u}_{E_C}]$, $\pi(u) = ((\underline{u}_{E_C} - u)pb + (u - pb)\pi(\underline{u}_{E_C})) / (\underline{u}_{E_C} - pb)$ and $(u, \pi(u))$ is supported by randomization between restricted empowerment and cooperative empowerment.*

(iii.) *For $u \in [\underline{u}_{E_C}, \bar{u}_{E_C}]$, $\pi(u) = \pi_{E_C}(u)$ and $(u, \pi(u))$ is supported by cooperative empowerment.*

(iv.) *For $u \in [\bar{u}_{E_C}, B]$, $\pi(u) = ((B - u)\pi(\bar{u}_{E_C}) + (u - \bar{u}_{E_C})b) / (B - \bar{u}_{E_C})$ and $(u, \pi(u))$ is supported by randomization between cooperative empowerment and unrestricted empowerment.*




Figure 2: This figure illustrates the feasible stage-game payoffs, the PPE payoff frontier, and the actions that support each point on the frontier. The dotted linear segments are supported by public randomization between their two endpoints, and this public randomization occurs at the end of the period.

We illustrate the lemma in Figure 2. The lemma shows that the payoff frontier is divided into four regions. In three of these four regions, payoffs are supported by randomization and, as a result, the payoff frontier is linear. In any such region, payoffs can be supported by multiple types of randomization. Since for all such randomizations, payoffs end up at one of the endpoints of the region eventually, we assume that the parties randomize between the endpoints immediately. In the remaining region, payoffs are supported by pure actions, and the payoff frontier is concave.

We can now describe the optimal relational contract and how it evolves over time.

PROPOSITION 1. *The optimal relational contract satisfies the following:*

First period: The agent's and the principal's payoffs are given by $u^ \in [\underline{u}_{E_C}, \bar{u}_{E_C}]$ and $\pi(u^*) = \pi_{E_C}(u^*)$. The parties engage in cooperative empowerment. If the agent chooses the principal's preferred project, his continuation payoff increases, and it falls otherwise.*

Subsequent periods: The agent's and the principal's expected payoffs are given by $u \in \{0\} \cup \{pb\} \cup [\underline{u}_{E_C}, \bar{u}_{E_C}] \cup \{B\}$ and $\pi(u)$. Their actions and continuation payoffs depend on what region u is in:

(i.) *If $u = 0$, the parties choose centralization. The agent's continuation payoff is given by $u_C(0) = 0$.*

(ii.) *If $u = pb$, the parties choose restricted empowerment. The agent's continuation payoff is given by $u_{E_R}(pb) = pb$.*

(iii.) If $u \in [\underline{u}_{EC}, \bar{u}_{EC}]$, the parties choose cooperative empowerment. If the agent chooses the principal's preferred project, his continuation payoff is given by $u_{EC,h}(u) > u$. If, instead, he chooses his own preferred project, his continuation payoff is given by $u_{EC,\ell}(u) < u$.

(iv.) If $u = B$, the parties engage in unrestricted empowerment. The agent's continuation payoff is given by $u_{EU}(B) = B$.

The proposition shows that the principal starts out by engaging in cooperative empowerment. To motivate the agent to choose her preferred project whenever it is available, the principal increases his continuation payoff whenever he chooses her preferred project, and she decreases his continuation payoff whenever he does not.

To see how the principal optimally increases the agent's continuation payoff, suppose the agent chooses the principal's preferred project for a number of consecutive periods. The principal then continues to engage in cooperative empowerment, and the agent's continuation payoff continues to increase, until the parties reach a period in which the continuation payoff passes the threshold \bar{u}_{EC} . At the end of that period, the parties engage in randomization to determine their actions in the following period. Depending on the outcome of this randomization, the principal either continues to engage in cooperative empowerment, or she moves to unrestricted empowerment. Finally, once play has moved to unrestricted empowerment, it remains there in all subsequent periods.

To see how the principal optimally decreases the agent's continuation payoff, suppose instead that the agent chooses his own preferred project for a number of consecutive periods. The principal then continues to engage in cooperative empowerment, and the agent's continuation payoff continues to decrease, until the parties reach a period in which the continuation payoff falls below the threshold \underline{u}_{EC} . At the end of that period, the parties engage in one of two types of randomization to determine their actions in the following period. If $u \in [pb, \underline{u}_{EC}]$, the principal either continues to engage in cooperative empowerment, or she moves to restricted empowerment. And if, instead, $u \in [0, pb)$, the principal either moves to restricted empowerment or she chooses centralization in the next period. Finally, once play has moved to either restricted empowerment or centralization, it remains there in all subsequent periods.

The above proposition leaves open two questions about the long-run outcome of the relationship. First, does the principal always end up administering a punishment or reward? And if she ever does administer a punishment, does it take the form of permanent centralization or permanent restricted empowerment? The next proposition answers these questions.

PROPOSITION 2. *In the optimal relational contract, the principal chooses cooperative empowerment for the first τ periods, where τ is random and finite with probability one. For $t > \tau$, the*

relationship results in unrestricted empowerment, restricted empowerment, or centralization forever. Both unrestricted empowerment and restricted empowerment are chosen with positive probability on the equilibrium path. If $B/b < (1 - \delta p) / (1 - \delta)$, centralization is never chosen on the equilibrium path. If $\bar{u}_{EC} < (1 - \delta) B + \delta pb$, centralization is chosen with positive probability on the equilibrium path.

The proposition shows that the answer to the first question—whether the principal always ends up administering a punishment or reward—is yes. And it shows that the answer to the second question—whether the punishment takes the form of permanent centralization or permanent restricted empowerment—is that it depends on the model’s parameters. Having characterized the optimal relational contract, we now turn to its implications.

The first implication is that ex ante identical firms can end up with long-run differences in their internal organizations which, in turn, create long-run differences in their performance levels. These differences arise solely because the firms experience different random events during their early histories. And they persist even though there are no informational or legal barriers that would prevent imitation. Instead, what prevents under-performing firms from imitating the organizations of their better-performing rivals is that their seemingly inefficient organizations are either a reward for past successes or a punishment for past failures. In either case, employees of under-performing firms would view the adoption of a different organizational structure as the violation of a mutual understanding and punish the firms accordingly. A firm’s history can therefore serve as a barrier to organizational imitation.

Specifically, the model predicts that firms either end up centralized—in which case the principal’s per period payoff is zero—or with some degree of empowerment—in which case the principal’s expected per period payoff is at least pB . As mentioned in the introduction, this prediction is consistent with Aghion et al. (2015), who find that decentralization is associated with better performance. More importantly, however, the model provides a potential explanation for why the centralized firms in their sample do not respond to the performance gap by becoming more decentralized, which is that they are constrained by their past promises.

The second implication is that the principal’s ability to make use of the agent’s information, and thus the payoff she is able to realize, declines over time. In particular, the principal’s first-period payoff $\pi(u^*)$ is strictly larger than the payoffs that the principal realizes once the relationship has converged to one of the steady states, in which case she gets at most b . Notice that the result that the principal’s payoff declines over time does not follow simply from the fact that we are focusing on an optimal relational contract that maximizes the principal’s equilibrium payoff. An optimal

relational contract could, in principle, require the parties to cycle between punishment and reward phases as it does in Padro i Miquel and Yared (2012), Li and Matouschek (2013), Zhu (2013), and Fong and Li (2015), who study games similar to ours, and as in the famous class of equilibria that Green and Porter (1984) focus on.

To see why such cycling is not optimal in our setting, notice that both rewards—letting the agent choose his preferred project even when the principal’s is available—and punishments—centralization or restricted empowerment—are costly for the principal. The threat to retract a previously promised reward, and the promise to retract a previously threatened punishment, however, do not impose any costs on the principal, yet they motivate the agent just the same. Delaying rewards and punishments therefore creates an additional and costless tool that the principal can use to motivate the agent.

6 The Failure to Exploit Public Opportunities

A key feature of our baseline model is that the changes in the environment that the organization fails to adapt to are privately observed by the agent. The model therefore cannot explain why some firms fail to take advantage of opportunities that can be identified without any special expertise and that are apparent to those at the top of the firms’ hierarchies and even to the wider public, as in our motivating examples.

To address this issue, we now divide the periods into a pre-opportunity phase and a post-opportunity phase. The only difference between the stage game in the post-opportunity phase and the one in our baseline model is that there are now two projects with known payoffs: the default project and the “new project.” If chosen and implemented, the new project yields (U_N, Π_N) net of implementation costs. The availability of the new project is publicly observable, and once the new project becomes available, it remains available in all future periods.




Figure 3: Timing of the stage game when public opportunities may become available.

The stage game in the pre-opportunity phase is similar to the one in our baseline model. The only difference is that at the end of the stage game in period t , just before the realization of the

public randomization device, nature determines whether the stage game in $t + 1$ will again be in the pre-opportunity phase or be the first in the post-opportunity phase. The probability that the game transitions to the post-opportunity phase is given by $q \in (0, 1)$, and whether this transition occurs is independent across periods. The timing is described in Figure 3.




Figure 4: This figure illustrates the PPE payoff frontier for the baseline model and the PPE payoff frontier and the actions that support each point on the frontier for the post-opportunity phase. The dotted segments are supported by public randomization between their endpoints, and this public randomization occurs at the end of the period. The principal's choice of action in the first period of the post-opportunity phase is the action associated with the point on the frontier associated with the continuation payoffs determined in the pre-opportunity phase.

To make the analysis interesting, we assume that the new project is neither too attractive nor too unattractive to both parties (see assumptions B1–B3 in appendix B). We denote by N the set (U_N, Π_N) that satisfy these conditions. In particular, these conditions ensure that $\Pi_N > B$, so that the new project is better for the principal than her preferred project in the baseline model. Notice that the game is now a stochastic game rather than a repeated one. To characterize the optimal relational contract, we therefore have to characterize two payoff frontiers: $\pi_{pre}(\cdot)$ —the frontier in the pre-opportunity phase—and $\pi_{post}(\cdot)$ —the frontier in the post-opportunity phase. Since the game transitions from the pre- to the post-opportunity phase, but never the reverse, we first characterize $\pi_{post}(\cdot)$ and then we partially characterize $\pi_{pre}(\cdot)$.

We characterize $\pi_{post}(\cdot)$ fully in the appendix. Figure 4 illustrates its main features and compares them to those of the payoff frontier $\pi(\cdot)$ in our baseline model. There are several differences to notice. First, because of the availability of the new opportunity, there are now two

additional arrangements that may be chosen in the the stage game. We refer to the arrangement in which the new opportunity is chosen and implemented with probability 1 as definite adoption (denoted by A_D). We refer to the arrangement in which the principal's project is chosen when available, and the new project is chosen otherwise as probabilistic adoption (denoted by A_P). Second, $\pi_{post}(\cdot)$ is everywhere above $\pi(\cdot)$. This reflects the fact that the principal's payoff from the new project Π_N is higher than her highest equilibrium payoff in the baseline model. In fact, Π_N is the highest payoff on the post-opportunity frontier. The principal would therefore always choose the new project if it were available in the first period. The final difference is that restricted empowerment is never chosen on the equilibrium path, since it is dominated by a randomization between definite adoption and centralization, both of which are equilibrium strategies of the stage game.

Now, consider the payoff frontier in the pre-opportunity phase. Recall that in the baseline model, payoffs on the frontier are supported by continuation payoffs that are again on the same frontier. In the pre-opportunity phase, in contrast, they are supported by continuation payoffs that are either on the frontier of the pre-opportunity phase or on the frontier of the post-opportunity phase. As the game evolves during the pre-opportunity phase, it can therefore become necessary to distort the continuation payoff that the agent receives if the new project becomes available away from U_N . It follows that there is then at least some chance that the principal will not choose the new project as soon as it becomes available and may, in fact, never do so. Proposition 3 provides conditions under which this is the case.

PROPOSITION 3. *For each $(U_N, \Pi_N) \in N$,*

- (i.) *There exists $\bar{\Pi}(U_N)$ and $\bar{q}(U_N, \Pi_N)$ such that for all $\Pi_N \leq \bar{\Pi}(U_N)$ and $q \leq \bar{q}(U_N, \Pi_N)$, there exists a public history h^T such that $\Pr(u_T = U_N | h^T) < 1$, where T is the first period in the post-opportunity phase.*
- (ii.) *There exists a $\hat{\delta}$ and $\hat{q}(U_N, \Pi_N)$ such that for all $\delta \leq \hat{\delta}$ and $q \leq \hat{q}(U_N, \Pi_N)$, there exists a public history h^T such that $\Pr(u_t = U_N | h^T) = 0$ for all $t \geq T$.*

The first part of the proposition provides conditions under which the principal does not choose the new project as soon as it becomes available. Suppose, for instance, that during the pre-opportunity phase the agent has chosen the principal's preferred project so often that his continuation payoff if the new project becomes available exceeds U_N . The principal then rewards the agent for his good performance in the pre-opportunity phase by promising not to choose the new project as soon as it becomes available.

The second part of the proposition shows that the principal may in fact promise never to

choose the new project. Suppose that during the pre-opportunity phase the agent has chosen the principal’s preferred project even more often so that his continuation payoff if the new project becomes available does not only exceed U_N but is actually equal to B . The principal is then rewarding the agent for his performance in the pre-opportunity phase by promising him that he will always be able to choose his own preferred project, even after the new project becomes available.

The same forces that limit the organization’s ability to adapt to changes in the environment that are privately observed by the agent can therefore also prevent it from taking advantage of opportunities, even when those opportunities are publicly observable and even when the principal can exploit those opportunities herself without having to induce the agent to do so.

7 Conclusions

Power is an inherently dynamic concept. The transfer of power from the top of an organization to those further down in the hierarchy is based on informal promises and is thus necessarily relational. Moreover, the allocation of power often evolves over time, with some members experiencing increases in their power while others see theirs slip away. To understand the allocation of power within organizations, this paper therefore develops a dynamic model of power.

In this model, the central purpose of power is to serve as a reward mechanism that those at the top use to discipline their subordinates and influence their decision-making. Even though this role of power as a reward has long been noted in sociology and organizational theory, it is absent in existing economic models of organizations. We capture this role formally and show that it gives rise to rich dynamics. Our results speak to how and why power is gained, lost, and retained and thus adds to our understanding of life inside organizations. And, in doing so, it provides new perspectives on why similar firms organize differently, even though those organizational differences are associated with differences in performance, and why established firms often have a harder time adapting to changes in their environments than their younger rivals, even when those changes are publicly observable. Finally, our results suggest that time matters for the choice between different organizational structures. This is in contrast to existing economic theories of organizations which, as we discussed in the introduction, are either static or focus on settings in which the optimal allocation of power is stationary. And, to our knowledge, it has not yet been explored in the emerging empirical literature on organizational design.

Our model is only a first step in developing a formal theory of power dynamics, and there are many issues that it does not address. Since we only allow for a single agent, for instance, we cannot explore horizontal differences in power across members of an organization, which are a

central concern in the sociology and organizational theory literature. We also take the boundaries of the firm as given and don't allow the principal to sell the formal authority to make decisions to the agent. A richer model would allow for such a transfer of formal authority and then develop a theory of the firm in which agents may integrate precisely because it allows them to use power dynamics as a reward mechanism. We leave these and related issues for future research.

References

- [1] AGHION, PHILIPPE, NICHOLAS BLOOM, AND JOHN VAN REENEN. 2013. Incomplete Contracts and the Internal Organization of Firms. *Journal of Law, Economics, and Organization*, 30(suppl 1): 37-63.
- [2] ———, NICHOLAS BLOOM, BRIAN LUCKING, RAFFAELLA SADUN, AND JOHN VAN REENEN. 2015. Growth and Decentralization in Bad Times. Mimeo.
- [3] ——— AND JEAN TIROLE. 1997. Formal and Real Authority in Organizations. *Journal of Political Economy*, 105(1): 1-29.
- [4] ABREU, DILIP, DAVID PEARCE, AND ENNIO STACCHETTI. 1990. Toward a Theory of Discounted Repeated Games with Imperfect Monitoring. *Econometrica*, 58(5): 1041-1063.
- [5] ALONSO, RICARDO AND NIKO MATOUSCHEK. 2007. Relational Delegation. *The RAND Journal of Economics*, 38(4): 1070-1089.
- [6] ATHEY, SUSAN AND KYLE BAGWELL. 2001. Optimal Collusion with Private Information. *The RAND Journal of Economics*, 32(3): 428-465.
- [7] BAKER, GEORGE, ROBERT GIBBONS, AND KEVIN MURPHY. 1999. Informal Authority in Organizations. *The Journal of Law, Economics, and Organization*, 15(1): 56-73.
- [8] BIAIS, BRUNO, THOMAS MARRIOTTI, GUILLAUME PLANTIN, AND JEAN-CHARLES ROCHET. 2007. Dynamic Security Design: Convergence to Continuous Time and Asset Pricing Implications. *The Review of Economic Studies*, 74(2): 345-390.
- [9] BLOOM, NICHOLAS, BENN EIFERT, APRAJIT MAHAJAN, DAVID MCKENZIE, AND JOHN ROBERTS. 2013. Does Management Matter? Evidence from India. *Quarterly Journal of Economics*, 128(1): 1-51.
- [10] BOLTON, PATRICK AND MATHIAS DEWATRIPONT. 2013. Authority in Organizations. In *Handbook of Organizational Economics*, eds. Robert Gibbons and John Roberts, Princeton University Press.
- [11] BOWER, JOSEPH AND CLAYTON CHRISTENSEN. 1995. Disruptive Technologies: Catching the Wave. *Harvard Business Review*, 73(1): 43-53.

- [12] CHASSANG, SYLVAIN. 2010. Building Routines: Learning, Cooperation, and the Dynamics of Incomplete Relational Contracts. *American Economic Review*, 100(1): 448-465.
- [13] CLEMENTI, GIAN LUCA AND HUGO HOPENHAYN. 2006. A Theory of Financing Constraints and Firm Dynamics. *The Quarterly Journal of Economics*, 121(1): 229-265.
- [14] CYERT, RICHARD AND JAMES MARCH. 1963. *A Behavioral Theory of the Firm*. Englewood Cliffs.
- [15] DEMARZO, PETER AND MICHAEL FISHMAN. 2007. Optimal Long-Term Financial Contracting. *The Review of Financial Studies*, 20(6): 2079-2128.
- [16] DESSEIN, WOUTER. 2002. Authority and Communication in Organizations. *The Review of Economic Studies*, 69(4): 811-838.
- [17] EICHENWALD, KURT. 2012. Microsoft's Lost Decade. *Vanity Fair*, 54(8).
- [18] EMERSON, RICHARD. 1962. Power-Dependence Relations. *American Sociological Review*, 27(1): 31-41.
- [19] FONG, YUK-FAI AND JIN LI. 2015. Relational Contracts, Efficiency Wages, and Employment Dynamics. Mimeo.
- [20] GARRETT, DANIEL AND ALESSANDRO PAVAN. 2012. Managerial Turnover in a Changing World. *Journal of Political Economy*, 120(5): 879-925.
- [21] GHOSH, PARIKSHIT AND DEBRAJ RAY. 1996. Cooperation in Community Interaction without Information Flows. *The Review of Economic Studies*, 63(3): 491-519.
- [22] GIBBONS, ROBERT AND REBECCA HENDERSON. 2013. What Do Managers Do? In *Handbook of Organizational Economics*, eds. Robert Gibbons and John Roberts, Princeton University Press.
- [23] GREEN, EDWARD AND ROBERT PORTER. 1984. Noncooperative Collusion under Imperfect Price Information. *Econometrica*, 52(1): 87-100.
- [24] GUO, YINGNI AND JOHANNES HORNER. 2015. Dynamic Mechanisms without Money. Mimeo.
- [25] HALAC, MARINA. 2014. Relationship Building: Conflict and Project Choice Over Time. *The Journal of Law, Economics, and Organization*, 30(4): 683-708.

- [26] ——— AND ANDREA PRAT. 2015. Managerial Attention and Worker Engagement. Mimeo.
- [27] HICKSON, DAVID, ROBERT HININGS, C.A. LEE, R.E. SCHNECK, AND JOHANNES PENNINGS. 1971. A Strategic Contingencies' Theory of Intraorganizational Power. *Administrative Science Quarterly*, 16(2): 216-229.
- [28] HOLMSTROM, BENGT. 1984. On the Theory of Delegation. In *Bayesian Models in Economic Theory*, eds. Marcel Boyer and Richard Kihlstrom. North-Holland.
- [29] KRANTON, RACHEL. 1996. The Formation of Cooperative Relationships. *Journal of Law, Economics, & Organization*, 12(1): 214-233.
- [30] LEVIN, JONATHAN. 2003. Relational Incentive Contracts. *American Economic Review*, 93(3): 835-857.
- [31] LI, JIN AND NIKO MATOUSCHEK. 2013. Managing Conflicts in Relational Contracts. *American Economic Review*, 103(6): 2328-2351.
- [32] LIPNOWSKI, ELLIOT AND JOAO RAMOS. 2015. Repeated Delegation. Mimeo.
- [33] MALCOMSON, JAMES. 1999. Individual Employment Contracts. In *Handbook of Labor Economics*, eds. Orley Ashenfelter and David Card, Vol. 3. Elsevier.
- [34] ———. 2013. Relational Incentive Contracts. In *Handbook of Organizational Economics*, eds. Robert Gibbons and John Roberts, Princeton University Press.
- [35] MAILATH, GEORGE AND LARRY SAMUELSON. 2001. Who Wants a Good Reputation? *The Review of Economic Studies*, 68(2): 415-441.
- [36] ——— AND ———. 2006. *Repeated Games and Reputations*. Oxford University Press.
- [37] MELUMAD, NAHUM AND TOSHIYUKI SHIBANO. 1991. Communication in Settings with No Transfers. *RAND Journal of Economics*, 22(2): 173-198.
- [38] MINTZBERG, HENRY. 1979. *The Structure of Organizations: A Synthesis of the Research*. Prentice Hall.
- [39] PADRO I MIQUEL, GERARD AND PIERRE YARED. 2012. The Political Economy of Indirect Control. *The Quarterly Journal of Economics*, 127(2): 947-1015.
- [40] PFEFFER, JEFFREY. 1981. *Power in Organizations*. Pitman.

- [41] ——— AND GERALD SALANCIK. 1978. *The External Control of Organizations: A Resource Dependence Perspective*. Harper & Row.
- [42] SAMUELSON, LARRY. 2006. The Economics of Relationships. *Econometric Society Monographs*, 41: 136–185.
- [43] SCHAEFER, SCOTT. 1998. Influence Costs, Structural Inertia and Organizational Change. *Journal of Economics and Management Strategy*, 7(2): 237–263.
- [44] SPEAR, STEPHEN AND SANJAY SRIVASTAVA. 1987. On Repeated Moral Hazard with Discounting. *Review of Economic Studies*, 54(4): 599–617.
- [45] WATSON, JOEL. 1999. Starting Small and Renegotiation. *Journal of Economic Theory*, 85(1): 52–90.
- [46] ———. 2002. Starting Small and Commitment. *Games and Economic Behavior*, 38(1): 176–199.
- [47] YANG, HUANXING. 2013. Non-stationary Relational Contracts with Adverse Selection. *International Economic Review*, 54(2): 525–547.
- [48] ZHU, JOHN. 2013. Optimal Contracts with Shirking. *The Review of Economic Studies*, 80(2): 812–839.

Appendix (For Online Publication)

This appendix is divided into two sections. Appendix A contains proofs for the results describing the optimal relational contract in the baseline model. Appendix B contains proofs for the model with public opportunities.

Appendix A: Optimal Relational Contract in the Baseline Model

LEMMA A1. *Without loss of generality, along the equilibrium path, $k_t = m_t$ for all t and $e_{it} = 1$ for $i = A, P$ for all t .*

Proof of Lemma A1. Take an equilibrium with $k_t \neq m_t$ for some t . Consider another strategy profile in which the agent's recommendation is k_t instead of m_t . This change does not affect any player's payoffs, so it does not affect any constraints. It follows that this new strategy profile is an equilibrium.

Consider any strategy profile in which for some t , $e_{At} = 1$ and $e_{Pt} = 0$, and consider an alternative strategy profile that coincides with the original strategy profile but for which $e_{At} = 0$. Under this strategy profile, the Principal's payoff is unaffected, the public outcome is unaffected, and the agent's payoff is strictly higher. This means that the original strategy profile cannot be an equilibrium. An identical argument shows that any equilibrium strategy profile cannot have $e_{Pt} = 1$ and $e_{At} = 0$ for any t . Therefore, $e_{At} = e_{Pt}$ for all t in any equilibrium.

Consider a strategy profile in which for some t , $e_{At} = e_{Pt} = 0$, and consider an alternative strategy profile that coincides with the original strategy profile but for which $k_t = D$ is chosen in that period. This change does not affect players' payoffs, and it does not affect any constraints, so it is also an equilibrium. ■

We now use the techniques developed by Abreu, Pearce, and Stacchetti (1990) to characterize the PPE payoff set and, in particular, its frontier. For this purpose, we define the payoff frontier as

$$\pi(u) \equiv \sup \{ \pi' : (u, \pi') \in \mathcal{E} \},$$

where \mathcal{E} is the PPE payoff set.

We can now state our first lemma, which establishes several properties of the PPE payoff set.

LEMMA A2. *The PPE payoff set \mathcal{E} has the following properties: (i.) it is compact; (ii.) $\pi(u)$ is concave; (iii.) $\inf\{u : (u, \pi) \in \mathcal{E}\} = 0$ and $\sup\{u : (u, \pi) \in \mathcal{E}\} = B$.*

Proof of Lemma A2: Part (i.): Note that there are finite number of actions the players can take, and standard arguments then imply that the PPE payoff set \mathcal{E} is compact. Part (ii.): the

concavity of π follows immediately from the availability of the public randomization device. Part (iii.): Notice that 0 is the agent's maxmin payoff. Moreover, $(0, 0)$ is an equilibrium payoff for the stage game, sustained by the strategy that players always choose the default project or, if they choose any other project, they both choose $e_i = 0$. It then follows that $\inf\{u : (u, \pi) \in \mathcal{E}\} = 0$. Also notice that B is the maximal feasible payoff for the agent. Moreover, (B, b) can be sustained as an equilibrium payoff in which the players choose entrenchment along the equilibrium path in every period. To see that this can be sustained as an equilibrium, notice that the agent does not have incentive to deviate since the equilibrium provides him with the highest feasible payoff. Any deviation by the principal would be an off-schedule deviation. The deviation can either be the choice of the default project, in which case the Principal would receive $0 < b$, or it can be the choice not to choose the agent's recommended project, in which case, it can be punished with both players choosing $e_i = 0$ in the implementation phase, in which case again, the Principal would receive $0 < b$. ■

LEMMA A3. *For any payoff $(u, \pi(u))$ on the frontier, the equilibrium continuation payoffs remain on the frontier.*

Proof of Lemma A3 : To show that for each payoff $(u, \pi(u))$ on the frontier, the equilibrium continuation payoffs remain on the frontier, it suffices to show that this is true if $(u, \pi(u))$ is supported by a pure action. Suppose $(u, \pi(u))$ is supported by centralization. Let (u_C, π_C) be the associated continuation payoff. Suppose to the contrary of the claim that $\pi_C < \pi(u_C)$. Now consider an alternative strategy profile that also specifies centralization but in which the continuation payoff is given by $(u_C, \hat{\pi}_C)$, where $\hat{\pi}_C = \pi_C + \varepsilon$ and where $\varepsilon > 0$ is small enough such that $\pi_C + \varepsilon \leq \pi(u_C)$. It follows from the promise-keeping constraints PK_C^P and PK_C^A that under this alternative strategy profile the payoffs are given by $\hat{u} = u$ and $\hat{\pi}_C = \pi(u) + \delta\varepsilon > \pi(u)$. It can be checked that this alternative strategy profile satisfies all the constraints and therefore constitutes a PPE. Since $\hat{\pi}_C > \pi(u)$, this contradicts the definition of $\pi(u)$, thus it must be that $\pi_C = \pi(u_C)$. The argument is identical when $(u, \pi(u))$ is supported by other actions. ■

LEMMA A4. *If any payoff pair $(u, \pi(u))$ is supported by a pure action, it is supported by an action $j \in \{C, E_R, E_C, E_U\}$.*

Proof of Lemma A4. It is without loss of generality to show that if a payoff pair (u, π) is supported by $m_t = A$ when P yields $(-\infty, \infty)$ and by $m_t = D$ when P yields (b, B) , then either $\pi < \pi(u)$ or there is a payoff-equivalent equilibrium in which the players randomize between choosing C or E_U . To see this, suppose that (u, π) is supported in this way. Define $u_1 = (1 - \delta)0 + \delta u_D$ and

$u_2 = (1 - \delta)B + \delta u_A$, where u_D is the continuation payoff associated with $k_t = D$ and u_A is the payoff associated with $k_t = A$. The agent's equilibrium utility is therefore

$$u = pu_1 + (1 - p)u_2,$$

and the principal's is

$$\begin{aligned}\pi &= p((1 - \delta)0 + \delta\pi(u_D)) + (1 - p)((1 - \delta)b + \delta\pi(u_A)) \\ &\leq p\pi(u_1) + (1 - p)\pi(u_2).\end{aligned}$$

Therefore, either $\pi < \pi(u)$ or (u, π) can be supported by randomization between $(u_1, \pi(u_1))$, where C is chosen in period t with probability 1, and $(u_2, \pi(u_2))$, where A is chosen in period t with probability 1. ■

For any action in $j \in \{C, E_R, E_C, E_U\}$, define $u_j(u)$ as the agent's continuation payoff when the equilibrium gives the agent payoff u . Note that for $j \in \{C, E_U\}$, the agent's continuation payoff is deterministic and is given by the corresponding promise-keeping constraints. The next lemma describes the agent's continuation payoff under restricted empowerment and cooperative empowerment.

LEMMA A5. The following hold.

(i.) If $(u, \pi(u))$ is supported by restricted empowerment, there exists a payoff-equivalent equilibrium in which the agent's continuation payoffs are $\delta u_{E_R, h}(u) = \delta u_{E_R, \ell}(u) = (u - (1 - \delta)pb) \equiv \delta u_{E_R}(u)$.

(ii.) If $(u, \pi(u))$ is supported by cooperative empowerment, there exists a payoff-equivalent equilibrium in which the agent's continuation payoffs are:

$$\begin{aligned}\delta u_{E_C, h}(u) &= u - (1 - \delta)b; \\ \delta u_{E_C, \ell}(u) &= u - (1 - \delta)B.\end{aligned}$$

Proof of Lemma A5: For part (i.), let $(u, \pi(u))$ be associated with the continuation payoffs $(u_{E_R, h}, \pi(u_{E_R, h}))$ and $(u_{E_R, \ell}, \pi(u_{E_R, \ell}))$. Suppose to the contrary that $u_{E_R, h} \neq u_{E_R, \ell}$. Consider an alternative strategy profile with continuation payoffs given by $(\hat{u}_{E_R, h}, \pi(\hat{u}_{E_R, h}))$ and $(\hat{u}_{E_R, \ell}, \pi(\hat{u}_{E_R, \ell}))$, where

$$\hat{u}_{E_R, h} = \hat{u}_{E_R, \ell} = pu_{E_R, h} + (1 - p)u_{E_R, \ell}.$$

Under this new strategy profile, $\text{PK}_{E_R}^A$ and IC_{E_R} still hold. This new profile gives the principal a payoff of

$$\begin{aligned}\hat{\pi} &= p[(1-\delta)B + \delta\pi(\hat{u}_{E_R,h})] + (1-p)\delta\pi(\hat{u}_{E_R,\ell}) \\ &\geq p[(1-\delta)B + \delta\pi(u_{E_R,h})] + (1-p)\delta\pi(u_{E_R,\ell}),\end{aligned}$$

where the inequality holds because π is concave. By $\text{PK}_{E_R}^A$, it then follows that $\delta u_{E_R,h} = \delta u_{E_R,\ell} = u - (1-\delta)pb$. We define this value to be δu_{E_R} .

For part (ii.), let $(u, \pi(u))$ be associated with the continuation payoffs $(u_{E_C,h}, \pi(u_{E_C,h}))$ and $(u_{E_C,\ell}, \pi(u_{E_C,\ell}))$. Suppose that for this PPE, IC_{E_C} is slack. That is, $(1-\delta)b + \delta u_{E_C,h} > (1-\delta)B + \delta u_{E_C,\ell}$. Now consider an alternative strategy profile with continuation payoffs given by $(\hat{u}_{E_C,h}, \pi(\hat{u}_{E_C,h}))$ and $(\hat{u}_{E_C,\ell}, \pi(\hat{u}_{E_C,\ell}))$, where $\hat{u}_{E_C,h} = u_{E_C,h} - (1-p)\varepsilon$ and $\hat{u}_{E_C,\ell} = u_{E_C,\ell} + p\varepsilon$ for $\varepsilon > 0$. It follows from the promise-keeping constraints $\text{PK}_{E_C}^P$ and $\text{PK}_{E_C}^A$ that, under this strategy profile, the payoffs are given by $\hat{u} = u$ and

$$\hat{\pi} = p[(1-\delta)B + \delta\pi(\hat{u}_{E_C,h})] + (1-p)[(1-\delta)b + \delta\pi(\hat{u}_{E_C,\ell})].$$

From the concavity of π it then follows that

$$\hat{\pi} \geq (1-\delta)b + \delta[(1-p)\pi(u_{E_C,\ell}) + p\pi(u_{E_C,h})] = \pi(u).$$

It can be checked that for sufficiently small ε this alternative strategy profile satisfies all the constraints and therefore constitutes a PPE. Since $\hat{\pi} \geq \pi(u)$ this implies that for any PPE with payoffs $(\pi, u(\pi))$ for which IC is not binding there exists another PPE for which IC_{E_C} is binding and which gives the parties weakly larger payoffs. Notice that when IC_{E_C} is binding, we have $u_{E_C,h}(u) = (u - (1-\delta)b)/\delta$ and $u_{E_C,\ell}(u) = (u - (1-\delta)B)/\delta$. This proves part (ii). ■

Next, let $\pi_j(u)$ for $j \in \{C, E_R, E_C, E_U\}$ be the principal's highest equilibrium payoff given action j and agent's payoff u . We then have

$$\begin{aligned}\pi_C(u) &= \delta\pi(u_C(u)), \\ \pi_{E_R}(u) &= p[(1-\delta)B] + \delta\pi(u_{E_R}(u)), \\ \pi_{E_C}(u) &= p[(1-\delta)B + \delta\pi(u_{E_C,h}(u))] + (1-p)[(1-\delta)b + \delta\pi(u_{E_C,\ell}(u))], \\ \pi_{E_U}(u) &= (1-\delta)b + \delta\pi(u_{E_U}(u)).\end{aligned}$$

LEMMA A6. *The PPE frontier $\pi(u)$ is the unique function that solves the following problem. For all $u \in [0, B]$,*

$$\pi(u) = \max_{\alpha_j \geq 0, u_j \in [0, B]} \sum_{j \in \{C, E_R, E_C, E_U\}} \alpha_j \pi_j(u_j)$$

such that

$$\sum_{j \in \{C, E_R, E_C, E_U\}} \alpha_j = 1$$

and

$$\sum_{j \in \{C, E_R, E_C, E_U\}} \alpha_j u_j = u.$$

Proof of Lemma A6: Since the frontier is Pareto efficient, by the APS bang-bang result, for any efficient payoff pair, only using the extreme points of the payoff set is sufficient. Replacing the sup with max is valid since the payoff set is compact. To establish the uniqueness, we just observe that the problem is now a maximization problem on a compact set, so that even if the maximizers are not unique, the maximum is. ■

LEMMA A7. *There exists a cutoff value $\bar{u}_{E_C} < B$ such that $\pi(u)$ is a straight line for $u \geq \bar{u}_{E_C}$ and $\pi_{E_U}(u) = \pi(u)$ if and only if $u \in [(1 - \delta)B + \delta\bar{u}_{E_C}, B]$.*

Proof of Lemma A7: First, notice that $\pi_{E_U}(B) = \pi(B)$. Next, recall that $\pi_{E_U}(u) = (1 - \delta)b + \delta\pi(u_{E_U}(u))$. Taking the right derivative, we have

$$\pi_{E_U}^+(u) = \delta\pi^+(u_{E_U}(u))u_{E_U}^+(u) = \pi^+(u_{E_U}(u)) \geq \pi^+(u),$$

where we used the fact that if $u < B$, $u_{E_U}(u) < u$, and therefore $\pi^+(u_{E_U}(u)) \geq \pi^+(u)$ by concavity of the frontier. Since $\pi_{E_U}^+(u) \geq \pi^+(u)$ for all $u < B$, there exists u^* such that $\pi_{E_U}(u) = \pi(u)$ if and only if $u \in [u^*, B]$.

Next, we show that $u^* < B$. That is, there exists some $u < B$ such that $\pi_{E_U}(u) = \pi(u)$. We prove this by contradiction. Suppose to the contrary that $\pi_{E_U}(u) < \pi(u)$ for all $u < B$. Choose a small enough $\varepsilon > 0$ such that $(B - \varepsilon, \pi(B - \varepsilon))$ cannot be supported by pure actions. Notice that such ε exists, because by assumption $(B - \varepsilon, \pi(B - \varepsilon))$ is not supported by E_U , and if it were supported by any other pure action, the agent's continuation payoff must exceed B , leading to a contradiction. This implies that $(B - \varepsilon, \pi(B - \varepsilon))$ must be supported by randomization, and therefore the frontier is a straight line between $B - \varepsilon$ and B . Denote the slope of the payoff frontier between $(B - \varepsilon, \pi(B - \varepsilon))$ and (B, b) as s . It then follows that for all $u \in [B - \delta\varepsilon, B]$ (i.e. $u_{E_U}(u) \geq B - \varepsilon$), we have

$$\pi_{E_U}(u) = \pi(u) = b + s(u - B).$$

This contradicts the assumption that $\pi_{E_U}(u) < \pi(u)$ for all $u < B$.

The above shows that $\pi_{E_U}(u) = \pi(u)$ for $u \in [u^*, B]$, where $u^* < B$. It follows that for all $u \in (u^*, B]$, $\pi_{E_U}^-(u) = \pi^-(u_{E_U}(u)) = \pi^-(u)$. Since π is concave, this implies that the slope of π is constant for all $u \in (u^*, B)$. That is, $\pi(u)$ is a straight line on $[u^*, B]$. Define $(\bar{u}_{E_C}, \pi(\bar{u}_{E_C}))$ to be the left endpoint of the line segment. ■

LEMMA A8. $\pi(u)$ is a straight line for $u \in [0, pb]$ and $\pi(u) = Bu/b$.

Proof of Lemma A8. Both $(0, 0)$ and (pb, pB) are stage-game equilibrium payoffs. Moreover, recall that the agent will never choose $e = 1$ for any project if the principal chooses $e = 0$. This implies that all payoffs fall weakly below the line that includes $(0, 0)$ and (pb, pB) . As a result, the line segment connecting $(0, 0)$ and (pb, pB) is on the frontier of the convex hull of the expected stage-game payoffs, which includes the PPE payoff set. ■

LEMMA A9. $\pi_C(u) = \pi(u)$ if and only if $u \in [0, \delta pb]$.

Proof of Lemma A9. First, note that $\pi_C(0) = \pi(0)$. Clearly, for all $u \in [0, \delta pb]$, $\pi_C(u) = \pi(u)$, because we have established that $\pi(u)$ is a straight line between $(0, 0)$ and (pb, pB) . If there exists $u > \delta pb$, then $u_C(u) > pb$. Then $\pi_C^-(u) = \pi^-(u_C(u)) < B/b$ since $\pi(u) < Bu/b$ for $u > pb$. If so, then for $\varepsilon > 0$ small enough, $\pi_C(pb - \varepsilon) > \pi(pb - \varepsilon)$, which is a contradiction. We therefore have that $\pi_C(u) = \pi(u)$ only in $[0, \delta pb]$. ■

LEMMA A10. There exists a cutoff value \underline{u}_{E_C} such that $\pi(u)$ is a straight line on $[pb, \underline{u}_{E_C}]$ and $\pi_{E_R}(u) = \pi(u)$ for all $u \in [(1 - \delta)pb, (1 - \delta)pb + \delta \underline{u}_{E_C}]$.

Proof of Lemma A10. We first argue that if $\pi_{E_R}(u) = \pi(u)$ for any $u > pb$, then π is linear on $[pb, u]$. We define \underline{u}_{E_C} as the right endpoint of this line segment and then show that for $u \in [pb, (1 - \delta)pb + \delta \underline{u}_{E_C}]$, $\pi_{E_R}(u) = \pi(u)$ and for any $u' > (1 - \delta)pb + \delta \underline{u}_{E_C}$, $\pi_{E_R}(u) < \pi(u)$. Finally, we show that $\pi_{E_R}(u) = \pi(u)$ for all $u \in [(1 - \delta)pb, pb]$.

For the first step, suppose $\pi(u) = \pi_{E_R}(u)$ for some $u > pb$. Then $u_{E_R}(u) > u$ and

$$\pi_{E_R}^-(u) = \pi^-(u_{E_R}(u)) \leq \pi^-(u)$$

since π is concave. This implies that for all $u' \in [pb, u]$, $\pi_{E_R}(u') \geq \pi(u')$ and therefore $\pi_{E_R}(u') = \pi(u')$. Moreover, it must be the case that $\pi^-(u_{E_R}(u)) = \pi^-(u)$, so π must be linear on $[pb, u]$. Define the right endpoint of this line segment as \underline{u}_{E_C} . For any $u \in [pb, \underline{u}_{E_C}]$ such that $u_R(u) \leq \underline{u}_{E_C}$, since π is linear between $[pb, \underline{u}_{E_C}]$, we can write $\pi(u) = pB + s(u - pb)$ for some s . Moreover,

$$\begin{aligned} \pi_{E_R}(u) &= (1 - \delta)pB + \delta\pi(u_{E_R}(u)) = (1 - \delta)pB + \delta(pB + s(u_{E_R}(u) - pb)) \\ &= pB + s(u - pb) = \pi(u). \end{aligned}$$

Next, suppose that $u_{E_R}(u) > \underline{u}_{E_C}$ and $\pi_{E_R}(u) = \pi(u)$. Then, since $u_{E_R}(u) > \underline{u}_{E_C}$,

$$\pi_{E_R}^-(u) = \pi^-(u_{E_R}(u)) < \pi^-(u).$$

Now, consider $\hat{u} = u - \varepsilon$ for ε small. Then $\pi_{E_R}(\hat{u}) > \pi(\hat{u})$, so it must be the case that $\pi_{E_R}(u) < \pi(u)$ for all u such that $u_{E_R}(u) > \underline{u}_{E_C}$.

Finally, since $\pi(u) = Bu/b$ on $[0, pb]$ and $0 \leq u_{E_R}(u) \leq pb$ whenever $(1 - \delta)pb \leq u \leq pb$,

$$\pi_{E_R}(u) = (1 - \delta)pB + \delta\pi(u_{E_R}(u)) = Bu/b = \pi(u).$$

This establishes that $\pi_{E_R}(u) = \pi(u)$ on $[(1 - \delta)pb, pb]$. ■

LEMMA A11. For all $u \in [\underline{u}_{E_C}, \bar{u}_{E_C}]$, $\pi_{E_C}(u) = \pi(u)$.

Proof of Lemma A11. By Lemmas A7, A9, and A10, for all $u \in [\underline{u}_{E_C}, \bar{u}_{E_C}]$, which is a subset of $[(1 - \delta)pb + \delta\underline{u}_{E_C}, (1 - \delta)B + \delta\bar{u}_{E_C}]$, if $(u, \pi(u))$ is supported by a pure action, then it must be supported by cooperative empowerment. Next, since \underline{u}_{E_C} and \bar{u}_{E_C} are extremal points, they must be supported by a pure action, and therefore $\pi_{E_C}(u) = \pi(u)$ for $u = \underline{u}_{E_C}$ and $u = \bar{u}_{E_C}$. Take any $u \in (\underline{u}_{E_C}, \bar{u}_{E_C})$. If $(u, \pi(u))$ is supported by randomization, it is supported by randomization between two points $(u_1, \pi(u_1))$ and $(u_2, \pi(u_2))$, $u_1 < u_2$, which are each supported by pure actions. If either $u_1 < \underline{u}_{E_C}$ or $u_2 > \bar{u}_{E_C}$, we can replace the left (right) endpoint of this randomization with $(\underline{u}_{E_C}, \pi(\underline{u}_{E_C}))$ ($(\bar{u}_{E_C}, \pi(\bar{u}_{E_C}))$), and this new randomization generates higher payoffs for the principal. Thus, if $(u, \pi(u))$ is supported by randomization, it is supported by randomization between two points that are each supported by cooperative empowerment.

Define the function $f(u) = \pi(u) - \pi_{E_C}(u)$ on $[\underline{u}_{E_C}, \bar{u}_{E_C}]$. $f(u)$ is continuous and therefore achieves a maximum on $[\underline{u}_{E_C}, \bar{u}_{E_C}]$. Suppose $f(u^*) = \max_{u \in [\underline{u}_{E_C}, \bar{u}_{E_C}]} f(u) > 0$. Then at $u = u^*$, $\pi(u^*) > \pi_{E_C}(u^*)$, and therefore $(u^*, \pi(u^*))$ is supported by randomization between two points $(u_1, \pi(u_1))$, $(u_2, \pi(u_2))$, each of which is supported by cooperative empowerment. But then $f(u^*) = \alpha f(u_1) + (1 - \alpha)f(u_2) = 0$, which implies that $\pi(u^*) = \pi_{E_C}(u^*)$. ■

LEMMA A12. $\bar{u}_{E_C} = (1 - \delta)b + \delta B$.

Proof or Lemma A12. Suppose $u_{E_C,h}(\bar{u}_{E_C}) < B$. Then, since by Lemma A11, $\pi_{E_C}(\bar{u}_{E_C}) = \pi(\bar{u}_{E_C})$, we have that $\pi^+(\bar{u}_{E_C}) = \pi_{E_C}^+(\bar{u}_{E_C}) = (1 - p)\pi^+(u_{E_C,\ell}(\bar{u}_{E_C})) + ps$, where s is the slope of the line segment between $(\bar{u}_{E_C}, \pi(\bar{u}_{E_C}))$ and (B, b) . Since $u_{E_C,\ell}(\bar{u}_{E_C}) < \bar{u}_{E_C}$, $\pi^+(u_{E_C,\ell}(\bar{u}_{E_C})) > s$. Take $\hat{u} = \bar{u}_{E_C} + \varepsilon$ for $\varepsilon > 0$ small. Then $\pi_{E_C}(\hat{u}) > \pi(\hat{u})$, which is a contradiction. Finally, since $u_{E_C,h}(\bar{u}_{E_C}) = B$, we have that $\bar{u}_{E_C} = (1 - \delta)b + \delta B$. ■

LEMMA A13. $\underline{u}_{E_C} \in [(1 - \delta)B, (1 - \delta)B + \delta pb]$.

Proof of Lemma A13. Suppose that $u_{E_C, \ell}(\underline{u}_{E_C}) > pb$. Then, since by Lemma A9, $\pi_{E_C}(\underline{u}_{E_C}) = \pi(\underline{u}_{E_C})$, we have that $\pi^-(\underline{u}_{E_C}) = \pi_{E_C}^-(\underline{u}_{E_C}) = (1-p)s + p\pi^-(u_{E_C, h}(\underline{u}_{E_C}))$, where s is the slope of the line segment between (pb, pB) and $(\underline{u}_{E_C}, \pi(\underline{u}_{E_C}))$. Since $u_{E_C, h}(\underline{u}_{E_C}) > \underline{u}_{E_C}$, $\pi^-(u_{E_C, h}(\underline{u}_{E_C})) < s$. Take $\hat{u} = \underline{u}_{E_C} - \varepsilon$ for $\varepsilon > 0$ small. Then $\pi_{E_C}(\hat{u}) > \pi(\hat{u})$, which is a contradiction. Since $u_{E_C, \ell}(\underline{u}_{E_C}) \in [0, pb]$, we have that $\underline{u}_{E_C} \in [(1-\delta)B, (1-\delta)B + \delta pb]$.

LEMMA A14. $b \leq \underline{u}_{E_C} \leq \max\{b, (1-\delta)B + \delta pb\}$.

Proof of Lemma A14. First, suppose that $\underline{u}_{E_C} < b$. Then we have $u_{E_C, \ell}(\underline{u}_{E_C}) < u_{E_C, h}(\underline{u}_{E_C}) < \underline{u}_{E_C}$. Since by Lemma A11, $\pi_{E_C}(\underline{u}_{E_C}) = \pi(\underline{u}_{E_C})$, we have that $\pi^+(\underline{u}_{E_C}) = \pi_{E_C}^+(\underline{u}_{E_C}) = (1-p)\pi^+(u_{E_C, \ell}(\underline{u}_{E_C})) + p\pi^+(u_{E_C, h}(\underline{u}_{E_C})) \geq s$, where s is the slope of the line segment between (pb, pB) and $(\underline{u}_{E_C}, \pi(\underline{u}_{E_C}))$. Take $\tilde{u} = \underline{u}_{E_C} + \varepsilon$ for $\varepsilon > 0$ small. It then follows that $\pi(\tilde{u}) \geq \pi_{E_C}(\tilde{u}) \geq \pi(u) + s\varepsilon$. This contradicts the definition of \underline{u}_{E_C} as the right end point of the line segment that includes (pb, pB) and $(\underline{u}_{E_C}, \pi(\underline{u}_{E_C}))$. This proves that $\underline{u}_{E_C} \geq b$.

Next, suppose $(1-\delta)B + \delta pb > b$ and suppose that $u_{E_C, \ell}(\underline{u}_{E_C}) > pb$. Then, since by Lemma A11, $\pi_{E_C}(\underline{u}_{E_C}) = \pi(\underline{u}_{E_C})$, we have that $\pi^-(\underline{u}_{E_C}) = \pi_{E_C}^-(\underline{u}_{E_C}) = (1-p)s + p\pi^-(u_{E_C, h}(\underline{u}_{E_C}))$, where recall again that s is the slope of the line segment between (pb, pB) and $(\underline{u}_{E_C}, \pi(\underline{u}_{E_C}))$. Since $u_{E_C, h}(\underline{u}_{E_C}) > \underline{u}_{E_C}$, which follows from the agent's promise-keeping constraint and that $(1-\delta)B + \delta pb > b$, we have $\pi^-(u_{E_C, h}(\underline{u}_{E_C})) < s$. Take $\hat{u} = \underline{u}_{E_C} - \varepsilon$ for $\varepsilon > 0$ small. Then $\pi_{E_C}(\hat{u}) > \pi(\hat{u})$, which is a contradiction. This proves that if $(1-\delta)B + \delta pb > b$, we have $u_{E_C, \ell}(\underline{u}_{E_C}) \leq pb$, and therefore, we have that $\underline{u}_{E_C} \leq (1-\delta)B + \delta pb$.

Finally, suppose $(1-\delta)B + \delta pb \leq b$, and suppose that $\underline{u}_{E_C} > b$. Then we have $u_{E_C, \ell}(\underline{u}_{E_C}) \in (pb, \underline{u}_{E_C})$ and $u_{E_C, h}(\underline{u}_{E_C}) > \underline{u}_{E_C}$. The same argument above implies that $\pi^-(\underline{u}_{E_C}) = \pi_{E_C}^-(\underline{u}_{E_C}) = (1-p)s + p\pi^-(u_{E_C, h}(\underline{u}_{E_C}))$. Again take $\hat{u} = \underline{u}_{E_C} - \varepsilon$ for $\varepsilon > 0$ small, we have $\pi_{E_C}(\hat{u}) > \pi(\hat{u})$, which is a contradiction. ■

PROPOSITION 1. *The optimal relational contract satisfies the following:*

First period: The agent's and the principal's payoffs are given by $u^ \in [\underline{u}_{E_C}, \bar{u}_{E_C}]$ and $\pi(u^*) = \pi_{E_C}(u^*)$. The parties engage in cooperative empowerment. If the agent chooses the principal's preferred project, his continuation payoff increases, and it falls otherwise.*

Subsequent periods: The agent's and the principal's expected payoffs are given by $u \in \{0\} \cup \{pb\} \cup [\underline{u}_{E_C}, \bar{u}_{E_C}] \cup \{B\}$ and $\pi(u)$. Their actions and continuation payoffs depend on what region u is in:

(i.) *If $u = 0$, the parties choose centralization. The agent's continuation payoff is given by $u_C(0) = 0$.*

(ii.) If $u = pb$, the parties choose restricted empowerment. The agent's continuation payoff is given by $u_{E_R}(pb) = pb$.

(iii.) If $u \in [\underline{u}_{E_C}, \bar{u}_{E_C}]$, the parties choose cooperative empowerment. If the agent chooses the principal's preferred project, his continuation payoff is given by $u_{E_C,h}(u) > u$. If, instead, he chooses his own preferred project, his continuation payoff is given by $u_{E_C,\ell}(u) < u$.

(iv.) If $u = B$, the parties engage in unrestricted empowerment. The agent's continuation payoff is given by $u_{E_U}(B) = B$.

Proof of Proposition 1. The preceding lemmas characterize the payoff frontier, the associated actions, and their continuation payoffs. It remains only to show that in the first period, parties engage in cooperative empowerment. Given our assumption that $b > pB$, it suffices to show that there exists an equilibrium payoff, sustained by cooperative empowerment, that gives the principal a payoff that exceeds b . In particular, consider $\pi(\bar{u}_{E_C})$, where recall that $\bar{u}_{E_C} = (1 - \delta)b + \delta B$. Notice that $u_{E_C,\ell}$ is decreasing in δ . It suffices to show that if $\pi(\bar{u}_{E_C}) \geq b$ when $u_{E_C,\ell}(\bar{u}_{E_C}) < pb$ for any $\hat{\delta}$, then $\pi(\bar{u}_{E_C}) \geq b$ for all $\delta \geq \hat{\delta}$. By Lemma A11,

$$\pi(\bar{u}_{E_C}) = p[(1 - \delta)B + \delta b] + (1 - p) \left[(1 - \delta)b + \delta \frac{B}{b} \left(B - \frac{1 - \delta}{\delta} (B - b) \right) \right]$$

It follows that

$$\frac{\pi(\bar{u}_{E_C}) - b}{b} = \frac{B - b}{b} \left[p(1 - \delta) + \delta(1 - p) + (2\delta - 1) \frac{B}{b} (1 - p) \right].$$

Notice that this expression is always positive if $2\delta \geq 1$. When $2\delta < 1$, Assumption (iii.) ensures that it is positive. ■

PROPOSITION 2. *In the optimal relational contract, the principal chooses cooperative empowerment for the first τ periods, where τ is random and finite with probability one. For $t > \tau$, the relationship results in unrestricted empowerment, restricted empowerment, or centralization forever. Both unrestricted empowerment and restricted empowerment are chosen with positive probability on the equilibrium path. Specifically, if $B/b < (1 - \delta p) / (1 - \delta)$, only restricted empowerment is chosen, and if $\bar{u}_{E_C} < (1 - \delta)B + \delta pb$, both restricted empowerment and centralization are chosen with positive probability.*

Proof of Proposition 2. Let $u^* = \arg\max_{u \in [0, B]} \pi(u)$ denote the agent's equilibrium utility when the principal's equilibrium utility is maximized. By Proposition 1, the relationship begins with cooperative empowerment, and therefore $u^* \geq \underline{u}_{E_C} \geq b$.

First, we will show that relationship settles in unrestricted empowerment with positive probability. To see this, first notice that $u^* > b$. Suppose to the contrary that $u^* = b$. Denote by s the

slope of the payoff frontier between (pb, pB) and $(b, \pi(b))$. Then

$$\pi^+(b) = \pi_{EC}^+(b) = (1-p)\pi^+(u_{EC,\ell}(b)) + p\pi^+(b).$$

As a result, $\pi^+(b) = \pi^+(u_{EC,\ell}(b)) \geq s > 0$, which contradicts the assumption that π is maximized at b . Next, given that $u^* > b$, we have that $u_h(u) - u > \frac{1-\delta}{\delta}(u^* - b)$ for all $u \in [u^*, \bar{u}_{EC}]$. Then there exists an $N > 0$ such that if the principal's preferred project is available in the first N periods, the agent's continuation payoff has to exceed \bar{u}_{EC} with probability, and therefore, with positive probability, the relationship settles in unrestricted empowerment.

Next, we provide conditions for which centralization is never chosen on the equilibrium path. Suppose $B/b < [1 - \delta p] / [1 - \delta]$. Then, $u_{EC,\ell}(b) > pb$, which means that for all $u \geq \underline{u}_{EC}$, $u_{EC,\ell}(b) > pb$. It follows that if u is ever below \underline{u}_{EC} , it will be above pb , and therefore centralization is reached with probability zero.

We now provide conditions for which centralization is chosen with positive probability on the equilibrium path. If $\bar{u}_{EC} < (1 - \delta)B + \delta pb$, then $u_{EC,\ell}(\bar{u}_{EC}) < pb$, which implies that wherever cooperative empowerment is used, the agent's continuation payoff falls below pb with positive probability, and therefore centralization is reached with positive probability.

Finally, by standard arguments, the agent's continuation payoff converges with probability one. ■

7.1 Appendix B: Optimal Relational Contract with Public Opportunities

Just as in the main section, we solve the game recursively by characterizing the PPE payoff sets. Define \mathcal{E}_{pre} as the PPE payoff set of the pre-opportunity phase and \mathcal{E}_{post} as the PPE payoff set of the post-opportunity phase. Let $\pi_i(u)$, $i \in \{pre, post\}$, be the associated payoff frontier. As in the baseline model, we can simplify our analysis by noting the following.

LEMMA B0. *Without loss of generality, along the equilibrium path, $k_t = m_t$ for all t and $e_{it} = 1$ for $i = A, P$ for all t .*

LEMMA B1. *For $i \in \{pre, post\}$, the PPE payoff set \mathcal{E}_i has the following properties: (i.) it is compact; (ii.) $\pi_i(u)$ is concave; (iii.) $\inf\{u : (u, \pi) \in \mathcal{E}\} = 0$ and $\sup\{u : (u, \pi) \in \mathcal{E}\} = B$.*

The proofs for these results are essentially the same as in the baseline model, and they are omitted here. Next, we list the actions that are used to sustain the equilibrium payoff set and the associated constraints.

Constraints in the Post-Opportunity Phase

We first list the set of constraints for supporting the PPE payoff set \mathcal{E}_{post} . Consider a PPE payoff pair $(u, \pi) \in \mathcal{E}_{post}$. As in the baseline model, we can restrict attention to the following arrangements: centralization, restricted empowerment, cooperative empowerment, unrestricted empowerment, opportunity, and strategic opportunity. As we will show, the optimal relational contract can be sustained without making use of any other arrangement. The first four arrangements are the same as in the baseline model.

Centralization Under centralization, the agent recommends the default project, and the principal chooses the default project. A payoff pair (u, π) can be supported by centralization if the following constraints are satisfied.

(i.) Feasibility: For the continuation payoffs to be feasible, they also need to be PPE payoffs. The continuation payoffs $u_{post,C}$ and $\pi_{post,C}$ that the parties realize under centralization therefore have to satisfy the self-enforcement constraint

$$(u_{post,C}, \pi_{post,C}) \in \mathcal{E}_{post}. \quad (\text{SE}_{post,C})$$

(ii.) No Deviation: As in the baseline model, the principal and the agent never want to deviate off schedule, and there are no feasible on-schedule deviations. In contrast to off-schedule deviations, on-schedule deviations are privately observed. Since the principal does not have any private information, and the agent does not get to choose a project, there are no on-schedule deviations under centralization.

(iii.) Promise Keeping: Finally, the consistency of the PPE payoff decomposition requires that the parties' payoffs are equal to the weighted sum of current and future payoffs. The promise-keeping constraints

$$\pi = \delta \pi_{post,C} \quad (\text{PK}_{post,C}^P)$$

and

$$u = \delta u_{post,C} \quad (\text{PK}_{post,C}^A)$$

ensure that this is the case.

Unrestricted Empowerment Under unrestricted empowerment, the agent always recommends his own preferred project, and the principal rubberstamps this recommendation. A payoff pair (u, π) can be supported by unrestricted empowerment if the following constraints are satisfied.

(i.) Feasibility: We denote by $(u_{post,E_U}, \pi_{post,E_U})$ the continuation payoffs under unrestricted empowerment. The self-enforcement constraint is then given by

$$(u_{post,E_U}, \pi_{post,E_U}) \in \mathcal{E}_{post}. \quad (\text{SE}_{post,E_U})$$

(ii.) No Deviation: As in the baseline model, the principal and the agent never want to deviate off schedule, and there are no feasible on-schedule deviations.

(iii.) Promise Keeping: The promise-keeping constraints are now given by

$$\pi = (1 - \delta) b + \delta \pi_{post,E_U} \quad (\text{PK}_{post,E_U}^P)$$

for the principal and

$$u = (1 - \delta) B + \delta u_{post,E_U} \quad (\text{PK}_{post,E_U}^A)$$

for the agent.

Cooperative Empowerment Under cooperative empowerment, the agent recommends the principal's preferred project when it is available and his own preferred project otherwise, and the principal rubberstamps the agent's recommendation. A payoff pair (u, π) can be supported by cooperative empowerment if the following constraints are satisfied.

(i.) Feasibility: For the continuation payoffs to be feasible, they also need to be PPE payoffs. Let $(u_{post,E_C,\ell}, \pi_{post,E_C,\ell})$ denote the parties' continuation payoffs if the agent chooses his own preferred project, and let $(u_{post,E_C,h}, \pi_{post,E_C,h})$ denote their payoffs if he chooses the principal's preferred project. The self-enforcement constraint is then given by

$$(u_{post,E_C,\ell}, \pi_{post,E_C,\ell}), (u_{post,E_C,h}, \pi_{post,E_C,h}) \in \mathcal{E}_{post}. \quad (\text{SE}_{post,E_C})$$

(ii.) No Deviation: The principal and the agent never want to deviate off schedule, and the principal has no on-schedule deviations. The agent, however, can deviate on schedule by recommending his preferred project when the principal's preferred project is available. The incentive constraint

$$(1 - \delta) b + \delta u_{post,E_C,h} \geq (1 - \delta) B + \delta u_{post,E_C,\ell} \quad (\text{IC}_{post,E_C})$$

ensures that he does not want to do so.

(iii.) Promise Keeping: The promise-keeping constraints are now given by

$$\pi = p[(1 - \delta) B + \delta \pi_{post,E_C,h}] + (1 - p)[(1 - \delta) b + \delta \pi_{post,E_C,\ell}] \quad (\text{PK}_{post,E_C}^P)$$

and

$$u = p[(1 - \delta) b + \delta u_{post,E_C,h}] + (1 - p)[(1 - \delta) B + \delta u_{post,E_C,\ell}]. \quad (\text{PK}_{post,E_C}^A)$$

Restricted Empowerment Under restricted empowerment, the agent recommends the principal's preferred project when it is available and the default project otherwise, and the principal always rubberstamps the agent's recommendation. A payoff pair (u, π) can be supported by restricted empowerment if the following constraints are satisfied.

(i.) Feasibility: For the continuation payoffs to be feasible, they also need to be PPE payoffs. Let $(u_{post,E_R,\ell}, \pi_{post,E_R,\ell})$ denote the parties' continuation payoffs if the agent recommends the default project, and let $(u_{post,E_R,h}, \pi_{post,E_R,h})$ denote their payoffs if he recommends the principal's preferred project. The self-enforcement constraint is then given by

$$(u_{post,E_R,\ell}, \pi_{post,E_R,\ell}), (u_{post,E_R,h}, \pi_{post,E_R,h}) \in \mathcal{E}_{post}. \quad (\text{SE}_{\text{post},E_R})$$

(ii.) No Deviation: The principal never wants to deviate off schedule. The agent can deviate off schedule by recommending his own project. If he does so, he receives $(1 - \delta)B$ this period followed by 0. To prevent the agent from deviating off schedule, we need that

$$u \geq (1 - \delta)B. \quad (\text{IC}_{\text{post},E_R}^{\text{Off}})$$

The agent can also deviate on schedule by recommending the default project when the principal's preferred project is available. The incentive constraint

$$(1 - \delta)b + \delta u_{post,E_R,h} \geq \delta u_{post,E_R,\ell} \quad (\text{IC}_{\text{post},E_R}^{\text{On}})$$

ensures that he does not want to do so.

(iii.) Promise Keeping: The promise-keeping constraints are now given by

$$\pi = p[(1 - \delta)B + \delta \pi_{post,E_R,h}] + (1 - p)\delta \pi_{post,E_R,\ell} \quad (\text{PK}_{\text{post},E_R}^{\text{P}})$$

and

$$u = p[(1 - \delta)b + \delta u_{post,E_R,h}] + (1 - p)\delta u_{post,E_R,\ell}. \quad (\text{PK}_{\text{post},E_R}^{\text{A}})$$

Definite Adoption Under definite adoption (A_D), the agent recommends the new project, and the principal chooses the new project. Note that a payoff pair (u, π) can be supported by definite adoption if the following constraints are satisfied.

(i.) Feasibility: We denote continuation payoffs under opportunity by $(u_{post,A_D}, \pi_{post,A_D})$. The self-enforcement constraint is then given by

$$(u_{post,A_D}, \pi_{post,A_D}) \in \mathcal{E}_{post}. \quad (\text{SE}_{\text{post},A_D})$$

(ii.) No Deviation: As in the case of centralization, the principal and the agent never wants to deviate off schedule, and there are no feasible on-schedule deviations.

(iii.) Promise Keeping: The promise-keeping constraints are now given by

$$\pi = (1 - \delta) \Pi_N + \delta \pi_{post, A_D} \quad (\text{PK}_{post, A_D}^P)$$

for the principal and

$$u = (1 - \delta) U_N + \delta u_{post, A_D} \quad (\text{PK}_{post, A_D}^A)$$

for the agent.

Probabilistic Adoption Under probabilistic adoption (A_P), the agent recommends the principal's preferred project when it is available and the new project otherwise, and the principal rubberstamps the agent's recommendation. A payoff pair (u, π) can be supported by probabilistic adoption if the following constraints are satisfied.

(i.) Feasibility: For the continuation payoffs to be feasible, they also need to be PPE payoffs. Let $(u_{post, A_P, \ell}, \pi_{post, A_P, \ell})$ denote the parties' continuation payoffs if the agent recommends the new project, and let $(u_{post, A_P, h}, \pi_{post, A_P, h})$ denote their payoffs if he recommends the principal's preferred project. The self-enforcement constraint is then given by

$$(u_{post, A_P, \ell}, \pi_{post, A_P, \ell}), (u_{post, A_P, h}, \pi_{post, A_P, h}) \in \mathcal{E}_{post}. \quad (\text{SE}_{post, A_P})$$

(ii.) No Deviation: The principal never wants to deviate off schedule. The agent can deviate off schedule by recommending his own project. If he does so, he receives $(1 - \delta) B$ this period followed by 0. To prevent the agent from deviating off-schedule, we need that

$$u \geq (1 - \delta) B. \quad (\text{IC}_{post, A_P}^{\text{Off}})$$

The agent can also deviate on schedule by recommending the new project when the principal's preferred project is available. The incentive constraint

$$(1 - \delta) b + \delta u_{post, A_P, h} \geq (1 - \delta) U_N + \delta u_{post, A_P, \ell} \quad (\text{IC}_{post, A_P}^{\text{On}})$$

ensures that he does not want to do so.

(iii.) Promise Keeping: The promise-keeping constraints are now given by

$$\pi = p[(1 - \delta) B + \delta \pi_{post, A_P, h}] + (1 - p)[(1 - \delta) \Pi_N + \delta \pi_{post, A_P, \ell}] \quad (\text{PK}_{post, A_P}^P)$$

and

$$u = p[(1 - \delta) b + \delta u_{post, A_P, h}] + (1 - p)[(1 - \delta) U_N + \delta u_{post, A_P, \ell}]. \quad (\text{PK}_{post, A_P}^A)$$

Randomization Finally, a payoff pair (u, π) can be supported by randomization. In this case, there exist at most three distinct PPE payoffs $(u_i, \pi_i) \in \mathcal{E}_{post}$, $i = 1, 2, 3$ such that

$$(u, \pi) = \alpha_1 (u_1, \pi_1) + \alpha_2 (u_2, \pi_2) + \alpha_3 (u_3, \pi_3)$$

for some $\alpha_1, \alpha_2, \alpha_3 \geq 0$ and $\alpha_1 + \alpha_2 + \alpha_3 = 1$.

Constraints in the Pre-Opportunity Phase

We now list the set of constraints for supporting the PPE payoff set \mathcal{E}_{pre} . Consider a PPE payoff pair $(u, \pi) \in \mathcal{E}_{pre}$. Again as in the baseline model, we can restrict our attention to the following arrangements: centralization, restricted empowerment, cooperative empowerment, and unrestricted empowerment. In contrast to the baseline model, we now need to specify the continuation payoffs both when the opportunity has arrived and when it has not.

Centralization Under centralization, the agent recommends the default project, and the principal chooses the default project. A payoff pair (u, π) can be supported by centralization if the following constraints are satisfied.

(i.) Feasibility: For the continuation payoffs to be feasible, they also need to be PPE payoffs. Let $(u_{pre,C}, \pi_{pre,C})$ be the associated continuation payoffs if the opportunity does not arrive next period and $(u_{trans,C}, \pi_{trans,C})$ be the associated continuation payoffs when the game transitions to the post-opportunity phase for the first time next period. The continuation payoffs therefore have to satisfy the self-enforcement constraint

$$(u_{pre,C}, \pi_{pre,C}) \in \mathcal{E}_{pre} \text{ and } (u_{trans,C}, \pi_{trans,C}) \in \mathcal{E}_{post} \quad (\text{SE}_{pre,C})$$

(ii.) No Deviation: As in the baseline model, the principal and the agent never want to deviate off schedule, and there are no feasible on-schedule deviations. Since the principal does not have any private information, and the agent does not get to choose a project, there are no on-schedule deviations under centralization.

(iii.) Promise Keeping: Finally, the consistency of the PPE payoff decomposition requires that the parties' payoffs are equal to the weighted sum of current and future payoffs. The promise-keeping constraints

$$\begin{aligned} \pi &= \delta [(1 - q) \pi_{pre,C} + q \pi_{trans,C}] ; & (\text{PK}_{pre,C}^P) \\ u &= \delta [(1 - q) u_{pre,C} + q u_{trans,C}] . & (\text{PK}_{pre,C}^A) \end{aligned}$$

Unrestricted Empowerment Under unrestricted empowerment, the agent always recommends his own preferred project, and the principal rubberstamps this recommendation. A payoff pair (u, π) can be supported by unrestricted empowerment if the following constraints are satisfied.

(i.) Feasibility: For the continuation payoffs to be feasible, they also need to be PPE payoffs. Let $(u_{pre,E_U}, \pi_{pre,E_U})$ be the associated continuation payoffs if the opportunity does not arrive next period and $(u_{trans,E_U}, \pi_{trans,E_U})$ be the associated continuation payoffs when the new opportunity arrives. The continuation payoffs therefore have to satisfy the self-enforcement constraint

$$(u_{pre,E_U}, \pi_{pre,E_U}) \in \mathcal{E}_{pre} \text{ and } (u_{trans,E_U}, \pi_{trans,E_U}) \in \mathcal{E}_{post} \quad (\text{SE}_{pre,E_U})$$

(ii.) No Deviation: As in the case of centralization, the principal and the agent never want to deviate off schedule, and there are no feasible on-schedule deviations.

(iii.) Promise Keeping: The promise-keeping constraints are now given by

$$\pi = (1 - \delta) b + \delta [(1 - q) \pi_{pre,E_U} + q \pi_{trans,E_U}] \quad (\text{PK}_{pre,E_U}^P)$$

for the principal and

$$u = (1 - \delta) B + \delta [(1 - q) u_{pre,E_U} + q u_{trans,E_U}] \quad (\text{PK}_{pre,E_U}^A)$$

for the agent.

Cooperative Empowerment Under cooperative empowerment, the agent recommends the principal's preferred project when it is available and his own preferred project otherwise, and the principal rubberstamps the agent's recommendation. A payoff pair (u, π) can be supported by cooperative empowerment if the following constraints are satisfied.

(i.) Feasibility: For the continuation payoffs to be feasible, they also need to be PPE payoffs. If the new opportunity does not arrive next period, let $(u_{pre,E_C,\ell}, \pi_{pre,E_C,\ell})$ denote the parties' continuation payoffs if the agent chooses his own preferred project, and $(u_{pre,E_C,h}, \pi_{pre,E_C,h})$ denote their payoffs if he chooses the principal's preferred project. Define $(u_{trans,E_C,\ell}, \pi_{trans,E_C,\ell})$ and $(u_{trans,E_C,h}, \pi_{trans,E_C,h})$ accordingly. The self-enforcement constraint is then given by

$$\begin{aligned} (u_{pre,E_C,\ell}, \pi_{pre,E_C,\ell}) &\in \mathcal{E}_{pre}, (u_{pre,E_C,h}, \pi_{pre,E_C,h}) \in \mathcal{E}_{pre}; & (\text{SE}_{pre,E_C}) \\ (u_{trans,E_C,\ell}, \pi_{trans,E_C,\ell}) &\in \mathcal{E}_{Post}, (u_{trans,E_C,h}, \pi_{trans,E_C,h}) \in \mathcal{E}_{Post}. \end{aligned}$$

(ii.) No Deviation: The principal and the agent never want to deviate off schedule, and the principal has no on-schedule deviations. The agent, however, can deviate on schedule by recommending his preferred project when the principal's preferred project is available. The incentive

constraint

$$(1 - \delta) b + \delta ((1 - q) u_{pre,EC,h} + q u_{trans,EC,h}) \geq (1 - \delta) B + \delta ((1 - q) u_{pre,EC,\ell} + q u_{trans,EC,\ell}). \quad (\text{IC}_{pre,EC})$$

ensures that he does not want to do so.

(iii.) Promise Keeping: The promise-keeping constraints are now given by

$$\begin{aligned} \pi &= p [(1 - \delta) B + \delta ((1 - q) \pi_{pre,EC,h} + q \pi_{trans,EC,h})] \\ &\quad + (1 - p) [(1 - \delta) b + \delta ((1 - q) \pi_{pre,EC,\ell} + q \pi_{trans,EC,\ell})], \end{aligned} \quad (\text{PK}_{pre,EC}^P)$$

and

$$\begin{aligned} u &= p [(1 - \delta) b + \delta ((1 - q) u_{pre,EC,h} + q u_{trans,EC,h})] \\ &\quad + (1 - p) [(1 - \delta) B + \delta ((1 - q) u_{pre,EC,\ell} + q u_{trans,EC,\ell})]. \end{aligned} \quad (\text{PK}_{pre,EC}^A)$$

Restricted Empowerment Under Restricted Empowerment, the agent recommends the principal's preferred project when it is available and the default project otherwise, and the principal always rubberstamps the agent's recommendation. A payoff pair (u, π) can be supported by restricted empowerment if the following constraints are satisfied.

(i.) Feasibility: For the continuation payoffs to be feasible, they also need to be PPE payoffs. If the new opportunity does not arrive next period, let $(u_{pre,ER,\ell}, \pi_{pre,ER,\ell})$ denote the parties' continuation payoffs if the agent chooses his own preferred project, and $(u_{pre,ER,h}, \pi_{pre,ER,h})$ denote their payoffs if he chooses the principal's preferred project. Define $(u_{trans,ER,\ell}, \pi_{trans,ER,\ell})$ and $(u_{trans,ER,h}, \pi_{trans,ER,h})$ accordingly. The self-enforcement constraint is then given by

$$\begin{aligned} (u_{pre,ER,\ell}, \pi_{pre,ER,\ell}) &\in \mathcal{E}_{pre}, (u_{pre,ER,h}, \pi_{pre,ER,h}) \in \mathcal{E}_{pre}; \\ (u_{trans,ER,\ell}, \pi_{trans,ER,\ell}) &\in \mathcal{E}_{Post}, (u_{trans,ER,h}, \pi_{trans,ER,h}) \in \mathcal{E}_{Post}. \end{aligned} \quad (\text{SE}_{pre,ER})$$

(ii.) No Deviation: The principal and the agent never want to deviate off schedule, and the principal has no on-schedule deviations. The agent, however, can deviate on schedule by recommending the default project when the principal's preferred project is available. The incentive constraint

$$(1 - \delta) b + \delta ((1 - q) u_{pre,ER,h} + q u_{trans,ER,h}) \geq \delta ((1 - q) u_{pre,ER,\ell} + q u_{trans,ER,\ell}) \quad (\text{IC}_{pre,ER})$$

ensures that he does not want to do so.

(iii.) **Promise Keeping:** The promise-keeping constraints are now given by

$$\begin{aligned}\pi &= p[(1-\delta)B + \delta((1-q)\pi_{pre,E_R,h} + q\pi_{trans,E_R,h})] \\ &\quad + (1-p)[(1-\delta)b + \delta((1-q)\pi_{pre,E_R,\ell} + q\pi_{trans,E_R,\ell})],\end{aligned}\tag{PK_{pre,E_R}^P}$$

and

$$\begin{aligned}u &= p[(1-\delta)b + \delta((1-q)u_{pre,E_R,h} + qu_{trans,E_R,h})] \\ &\quad + (1-p)[(1-\delta)B + \delta((1-q)u_{pre,E_R,\ell} + qu_{trans,E_R,\ell})].\end{aligned}\tag{PK_{pre,E_R}^A}$$

Randomization Finally, a payoff pair (u, π) can be supported by randomization. In this case, there exist at most three distinct PPE payoffs $(u_i, \pi_i) \in \mathcal{E}_{Pre}$, $i = 1, 2, 3$ such that

$$(u, \pi) = \alpha_1(u_1, \pi_1) + \alpha_2(u_2, \pi_2) + \alpha_3(u_3, \pi_3)$$

for some $\alpha_1, \alpha_2, \alpha_3 \geq 0$ and $\alpha_1 + \alpha_2 + \alpha_3 = 1$.

Properties of π_{post}

To focus our analysis but to allow for sufficient generality, we make the following assumptions.

ASSUMPTION B1. $pb < U_N \leq b$.

ASSUMPTION B2. $(1-\delta)B \leq pb + (1-p)U_N$.

ASSUMPTION B3. $B < \Pi_N \leq \min\{B - U_N, (p^{-1} + 1 - p)B\}$.

We will refer to the set (U_N, Π_N) that satisfy assumptions B1, B2, and B3 as N . Lemma B2 shows that $\pi_{post}(u)$ shares similar features as the PPE payoff frontier in the main section.

LEMMA B2. *For any payoff $(u, \pi_{Post}(u))$ on the frontier, the equilibrium continuation payoffs remain on the frontier. For all $(U_N, \Pi_N) \in N$, the following hold.*

(i.) *If $(u, \pi_{post}(u))$ is supported with centralization, the agent's continuation payoff is given by*

$$\delta u_{post,C}(u) = u.$$

(ii.) *If $(u, \pi_{post}(u))$ is supported with unrestricted empowerment, the agent's continuation payoff is given by*

$$\delta u_{post,E_U}(u) = u - (1-\delta)B.$$

(iii.) If $(u, \pi_{post}(u))$ is supported with cooperative empowerment, the agent's continuation payoff can be chosen to be

$$\begin{aligned}\delta u_{post, E_C, h}(u) &= u - (1 - \delta) b; \\ \delta u_{post, E_C, \ell}(u) &= u - (1 - \delta) B.\end{aligned}$$

(iv.) If $(u, \pi_{post}(u))$ is supported with restricted empowerment, the agent's continuation payoff is given by

$$\delta u_{post, E_R, h}(u) = \delta u_{post, E_R, \ell}(u) = u - (1 - \delta) pb.$$

(v.) If $(u, \pi_{post}(u))$ is supported with definite adoption, the agent's continuation payoff is given by

$$\delta u_{post, A_D}(u) = u - (1 - \delta) U_N.$$

(vi.) If $(u, \pi_{post}(u))$ is supported with probabilistic adoption, the agent's continuation payoff is given by

$$\delta u_{post, A_P, h}(u) = \delta u_{post, A_P, \ell}(u) = u - (1 - \delta) (pb + (1 - p) U_N).$$

Proof of Lemma B2: Parts (i.)–(iv.) are proven in the same way as in the proof of the baseline model. Part (v.) follows directly from the agent's promise-keeping condition (PK_{post, A_D}^A) . Part (vi.) follows from the agent's promise-keeping condition (PK_{post, A_P}^A) and the condition that $b \geq u_N$, which ensures that the agent's on-schedule IC constraint is satisfied when $\delta u_{post, A_P, h}(u) = \delta u_{post, A_P, \ell}(u)$. ■

Just as in the main section, let $\pi_{post, j}(u)$ for $j \in \{C, E_R, E_C, E_U, A_D, A_P\}$ be the highest equilibrium payoff for the principal when the agent's payoff is u and action j is chosen. Lemma B2 implies that

$$\begin{aligned}\pi_{post, C}(u) &= \delta \pi_{post}(u_{post, C}(u)); \\ \pi_{post, E_R}(u) &= (1 - \delta) pB + \delta \pi_{post}(u_{post, E_R}(u)); \\ \pi_{post, E_C}(u) &= p[(1 - \delta) B + \delta \pi_{post}(u_{post, E_C, h}(u))] + (1 - p)[(1 - \delta) b + \delta \pi_{post}(u_{post, E_C, \ell}(u))]; \\ \pi_{post, E_U}(u) &= (1 - \delta) b + \delta \pi_{post}(u_{post, E_U}(u)); \\ \pi_{post, A_D}(u) &= (1 - \delta) \Pi_N + \delta \pi_{post}(u_{post, A_D}(u)); \\ \pi_{post, A_P}(u) &= (1 - \delta) (pB + (1 - p) \Pi_N) + \delta \pi_{post}(u_{post, A_P}(u)).\end{aligned}$$

The characterization of π_{post} is similar to the analysis in the baseline model. It is worth noting that if $(U_N, \Pi_N) \in N$, restrictive empowerment is no longer used to support any payoff pair $(u, \pi_{post}(u))$.

LEMMA B3. *For each $(U_N, \Pi_N) \in N$, there exist two cutoffs $\underline{u}_{post, EC}$ and $\bar{u}_{post, EC}$ such that the PPE payoff frontier $\pi_{post}(u)$ is divided into at most five regions:*

(i.) *For $u \in (0, U_N)$, $\pi_{post}(u)$ is supported by randomization between centralization and definite adoption. $\pi_{post}(0) = 0$ and $\pi_{post}(U_N) = \Pi_N$.*

(ii.) *For $u \in (U_N, pb + (1-p)U_N]$, $\pi_{post}(u)$ is supported by randomization between definite adoption and probabilistic adoption. $\pi_{post}(pb + (1-p)U_N) = pB + (1-p)\Pi_N$.*

(iii.) *For $u \in [pb + (1-p)U_N, \underline{u}_{post, EC}]$, $\pi_{post}(u)$ is supported by randomization between probabilistic adoption and cooperative empowerment.*

(iv.) *For $u \in [\underline{u}_{post, EC}, \bar{u}_{post, EC}]$, $\pi_{post}(u)$ is supported by cooperative empowerment.*

(v.) *For $u \in [\bar{u}_{post, EC}, B]$, $\pi_{post}(u)$ is supported by randomization between cooperative empowerment and unrestricted empowerment.*

In addition, $\bar{u}_{post, EC} = (1-\delta)b + \delta B$; $b \leq \underline{u}_{post, EC} \leq \max\{b, (1-\delta)B + \delta(pb + (1-p)U_N)\}$. The payoff frontier π_{post} is maximized at U_N .

Proof of Lemma B3: To see part (i.), note that $(0, 0)$ and (U_N, Π_N) are stage-game equilibrium payoffs. Recall that the agent will never choose $e = 1$ for any project if the principal chooses $e = 0$. This implies that all equilibrium payoffs lie weakly below the line segment that connects $(0, 0)$ and (U_N, Π_N) . As a result, the line segment connecting $(0, 0)$ and (U_N, Π_N) is on the frontier of the convex hull of the expected stage-game payoffs, which includes the PPE payoff set. For part (ii.), notice that $(pb + (1-p)U_N, pB + (1-p)\Pi_N)$ is a stage-game equilibrium expected payoff given that $(1-\delta)B \geq pb + (1-p)U_N$. Notice that $(pb + (1-p)U_N, pB + (1-p)\Pi_N)$ is on the line segment between (U_N, Π_N) and (b, B) . This line segment is on the frontier of the convex hull of the expected stage-game payoffs, which includes the PPE payoff set. For the remaining part of the lemma, notice that for the proof of parts (iii.)–(v.), the value of $\bar{u}_{post, EC}$ and the bounds on $\underline{u}_{post, EC}$ follow from the same analysis as in the baseline model. Finally, since (U_N, Π_N) is an equilibrium payoff, and Π_N is the highest stage-game payoff for the principal, it is immediate that π_{post} is maximized at U_N . ■

Properties of π_{pre}

Now we characterize the payoff frontier of the pre-opportunity game. Unlike the analysis of the baseline model or of the post-opportunity game, there are no explicit expressions for the agent's

continuation payoffs. Instead, they are pinned down by the following two conditions. First, their expected value is determined by the promise-keeping condition (with the same expressions as those in the baseline model). Second, we have $\pi'_{pre}(u_{pre,j}(u)) = \pi'_{post}(u_{trans,j}(u))$ for $j = \{C, E_R, E_U, (E_C, h), (E_C, \ell)\}$ when the payoff frontiers are differentiable. The next lemma provides the details.

LEMMA B4. *For any payoff $(u, \pi_{pre}(u))$ on the frontier, the equilibrium continuation payoffs remain on the frontier. In addition, the following holds.*

(i.) *If $(u, \pi_{pre}(u))$ is supported by centralization, the agent's continuation payoff satisfies*

$$\delta q u_{trans,C}(u) + \delta(1-q) u_{pre,C}(u) = u.$$

In addition,

$$\pi_{pre}^+(u_{pre,C}(u)) \leq \pi_{post}^-(u_{trans,C}(u)); \quad \pi_{post}^+(u_{trans,C}(u)) \leq \pi_{pre}^-(u_{pre,C}(u)).$$

(ii.) *If $(u, \pi_{pre}(u))$ is supported by restricted empowerment, the agent's continuation payoff satisfies $u_{pre,E_R,\ell}(u) = u_{pre,E_R,h}(u) \equiv u_{pre,E_R}(u)$, $u_{trans,E_R,\ell}(u) = u_{trans,E_R,h}(u) \equiv u_{trans,E_R}(u)$*

$$\delta [q u_{trans,E_R}(u) + (1-q) u_{pre,E_R}(u)] = u - (1-\delta)pb.$$

In addition,

$$\pi_{pre}^+(u_{pre,E_R}(u)) \leq \pi_{post}^-(u_{trans,E_R}(u)); \quad \pi_{post}^+(u_{trans,E_R}(u)) \leq \pi_{pre}^-(u_{pre,E_R}(u)).$$

(iii.) *If $(u, \pi_{pre}(u))$ is supported by cooperative empowerment, the agent's continuation payoff can be chosen to satisfy*

$$\begin{aligned} \delta q u_{trans,E_C,\ell}(u) + \delta(1-q) u_{pre,E_C,\ell}(u) &= u - (1-\delta)B; \\ \delta q u_{trans,E_C,h}(u) + \delta(1-q) u_{pre,E_C,h}(u) &= u - (1-\delta)b. \end{aligned}$$

In addition, for $j \in \{h, \ell\}$,

$$\pi_{pre}^+(u_{pre,E_C,j}(u)) \leq \pi_{post}^-(u_{trans,E_C,j}(u)); \quad \pi_{post}^+(u_{trans,E_C,j}(u)) \leq \pi_{pre}^-(u_{pre,E_C,j}(u)).$$

(iv.) *If $(u, \pi_{pre}(u))$ is supported by unrestricted empowerment, the agent's continuation payoff is given by*

$$\delta q u_{trans,E_U}(u) + \delta(1-q) u_{trans,E_U}(u) = u - (1-\delta)B.$$

In addition,

$$\pi_{pre}^+(u_{pre,E_U}(u)) \leq \pi_{post}^-(u_{trans,E_U}(u)); \quad \pi_{post}^+(u_{trans,E_U}(u)) \leq \pi_{pre}^-(u_{pre,E_U}(u)).$$

Proof of Lemma B4: This is proven in the same way as that in the baseline model. The additional inequality constraints arise, because at the optimum, for a given expected continuation payoff for the agent, it has to be optimal for the principal not to increase or decrease the agent's state-contingent continuation payoff. ■

Now we can prove proposition 3.

PROPOSITION 3. *For each $(U_N, \Pi_N) \in N$,*

- (i.) *There exists $\bar{\Pi}(U_N)$ and $\bar{q}(U_N, \Pi_N)$ such that for all $\Pi_N \leq \bar{\Pi}(U_N)$ and $q \leq \bar{q}(U_N, \Pi_N)$, there exists a public history h^T such that $\Pr(u_T = U_N | h^T) < 1$, where T is the first period in the post-opportunity phase.*
- (ii.) *There exists a $\hat{\delta}$ and $\hat{q}(U_N, \Pi_N)$ such that for all $\delta \leq \hat{\delta}$ and $q \leq \hat{q}(U_N, \Pi_N)$, there exists a public history h^T such that $\Pr(u_t = U_N | h^T) = 0$ for all $t \geq T$.*

Proof of Proposition 3: Denote $\pi_{pre}^q(u)$ to be the payoff frontier in the pre-opportunity game with parameter q , and notice that $\pi_{pre}^0(u) = \pi(u)$, which is the frontier of the baseline model. By Berge's maximum theorem, $\lim_{q \rightarrow 0} \pi_{pre}^q(u) = \pi(u)$ for each u . Define $\bar{u}_{pre, EC}^q = \max \left\{ u : \pi_{pre}^q(u) = \pi_{pre, EC}^q(u) \right\}$. Then, $\lim_{q \rightarrow 0} \bar{u}_{pre, EC}^q = \bar{u}_{pre, EC}^0 = \bar{u}_{EC}$ and $\pi_{pre}^q(u)$ is sustained by randomization on the interval $(\bar{u}_{pre, EC}^q, \tilde{u}_{pre, EC}^q)$ for some $\tilde{u}_{pre, EC}^q > \bar{u}_{pre, EC}^q$. Denote s^q to be the slope of π_{pre}^q on this interval, and denote by s^0 the slope of π on this interval. It follows that $\lim_{q \rightarrow 0} s^q = s^0$.

To prove part (i.), it suffices to show that we cannot have both $u_{trans, EC, h}^q(\bar{u}_{pre, EC}^q) = U_N$ and $u_{trans, EC, \ell}^q(\bar{u}_{pre, EC}^q) = U_N$. In order to get a contradiction, suppose to the contrary that $(\bar{u}_{pre, EC}^q, \pi_{pre}^q(\bar{u}_{pre, EC}^q))$ is supported by cooperative empowerment and the continuation payoffs $(u_{trans, EC, h}^q(\bar{u}_{pre, EC}^q), u_{trans, EC, \ell}^q(\bar{u}_{pre, EC}^q))$ and $(u_{pre, EC, h}^q(\bar{u}_{pre, EC}^q), u_{pre, EC, \ell}^q(\bar{u}_{pre, EC}^q))$, where $u_{trans, EC, h}^q(\bar{u}_{pre, EC}^q) = u_{trans, EC, \ell}^q(\bar{u}_{pre, EC}^q) = U_N$. Consider an alternative strategy profile that delivers equilibrium payoffs $(\hat{u}, \hat{\pi})$ on the frontier, and this point is sustained by cooperative empowerment with continuation payoffs given by $u_{trans, EC, h} = u_{trans, EC, \ell} = U_N + \varepsilon$ for some $\varepsilon > 0$ small and $u_{pre, EC, h} = u_{pre, EC, h}^q(\bar{u}_{pre, EC}^q)$ and $u_{pre, EC, \ell} = u_{pre, EC, \ell}^q(\bar{u}_{pre, EC}^q)$. The promise-keeping condition implies that

$$\hat{u} = \bar{u}_{pre, EC}^q + \delta \varepsilon q$$

and, if we denote by r the slope between (U_N, Π_N) and $(pb + (1-p)U_N, pB + (1-p)\Pi_N)$,

$$\hat{\pi} = \pi_{pre}^q(\bar{u}_{pre, EC}^q) + \delta q r \varepsilon.$$

Now, for any $U_N \leq b$, there exists $\bar{\Pi}(U_N)$ such that for all $B < \Pi_N \leq \bar{\Pi}(u_N)$, the slope $r > s^0/2$. Further, there exists $\bar{q}(U_N)$ such that for any $q \leq \bar{q}(U_N)$, $s^q \in (3s^0/4, s^0)$. It then follows that

$$\hat{\pi} > \pi_{pre}^q(\bar{u}_{pre,EC}^q) + \frac{1}{2}\delta qs^0\varepsilon$$

and

$$\begin{aligned} \pi_{pre}^q(\bar{u}_{pre,EC}^q + \delta q\varepsilon) &= \pi_{pre}^q(\bar{u}_{pre,EC}^q) + \delta qs^q\varepsilon \\ &\leq \pi_{pre}^q(\bar{u}_{pre,EC}^q) + \frac{3}{4}\delta qs^0\varepsilon \\ &< \pi_{pre}^q(\bar{u}_{pre,EC}^q) + \frac{1}{2}\delta qs^0\varepsilon < \hat{\pi}, \end{aligned}$$

which implies that $(\hat{u}, \hat{\pi})$ lies above the point $(\hat{u}, \pi_{pre}^q(\hat{u}))$ because $s^0 < 0$, which is a contradiction.

To prove part (ii.), suppose that $\delta < \hat{\delta} = \frac{B-b}{2B-(1+p)b}$, so that $u_{EC,\ell}(\bar{u}_{EC}) < pb$. It suffices to show that for q sufficiently small, $u_{pre,EC,h}(\bar{u}_{pre,EC}^q) = B$. In order to get a contradiction, suppose that $u_{pre,EC,h}(\bar{u}_{pre,EC}^q) < B$ for all q . Define s^q as above. We know that $\lim_{q \rightarrow 0} s^q = s^0$. As above, suppose to the contrary that $(\bar{u}_{pre,EC}^q, \pi_{pre}^q(\bar{u}_{pre,EC}^q))$ is supported by cooperative empowerment and continuation payoffs $(u_{trans,EC,h}^q(\bar{u}_{pre,EC}^q), u_{trans,EC,\ell}^q(\bar{u}_{pre,EC}^q))$ and $(u_{pre,EC,h}^q(\bar{u}_{pre,EC}^q), u_{pre,EC,\ell}^q(\bar{u}_{pre,EC}^q))$, where $u_{pre,EC,h}^q(\bar{u}_{pre,EC}^q) < B$. Consider an alternative strategy profile that delivers equilibrium payoffs $(\hat{u}, \hat{\pi})$ on the frontier, and this point is sustained by cooperative empowerment with continuation payoffs given by $\hat{u}_{trans,EC,h}(\bar{u}_{pre,EC}^q) = u_{trans,EC,h}^q(\bar{u}_{pre,EC}^q)$ and $\hat{u}_{trans,EC,\ell}(\bar{u}_{pre,EC}^q) = u_{trans,EC,\ell}^q(\bar{u}_{pre,EC}^q)$, for the transitional continuation payoffs, and $\hat{u}_{pre,EC,h}(\bar{u}_{pre,EC}^q) = u_{pre,EC,h}^q(\bar{u}_{pre,EC}^q) + \varepsilon$ and $\hat{u}_{pre,EC,\ell}(\bar{u}_{pre,EC}^q) = u_{pre,EC,\ell}^q(\bar{u}_{pre,EC}^q) + \varepsilon$ for the continuation payoffs that remain on the pre-opportunity frontier. This new strategy profile provides the agent with a payoff of

$$\hat{u} = \bar{u}_{pre,EC}^q + \delta(1-q)\varepsilon$$

and the principal with a payoff of

$$\begin{aligned} \hat{\pi} &= \pi_{pre}^q(\bar{u}_{pre,EC}^q) + \delta[p(1-q)(\pi_{pre}^q(\hat{u}_{pre,EC,h}) - \pi_{pre}^q(u_{pre,EC,h}))] \\ &\quad + \delta(1-p)(1-q)(\pi_{pre}^q(\hat{u}_{pre,EC,\ell}) - \pi_{pre}^q(u_{pre,EC,\ell})). \end{aligned}$$

Moreover, this change preserves the agent's incentive constraint, and it is an equilibrium payoff.

Next, notice that

$$\pi_{pre}^q(\hat{u}) = \pi_{pre}^q(u) + s^q\delta(1-q)\varepsilon$$

and

$$\hat{\pi} \geq \pi_{pre}^q(u) + \delta(1-q)\varepsilon \left[p\pi_{pre}^{q+} \left(u_{pre,E_C,h} \left(\bar{u}_{pre,E_C}^q \right) \right) + (1-p)\pi_{pre}^{q+} \left(u_{pre,E_C,\ell} \left(\bar{u}_{pre,E_C}^q \right) \right) \right].$$

Therefore, we obtain a contradiction if

$$p\pi_{pre}^{q+} \left(u_{pre,E_C,h} \left(\bar{u}_{pre,E_C}^q \right) \right) + (1-p)\pi_{pre}^{q+} \left(u_{pre,E_C,\ell} \left(\bar{u}_{pre,E_C}^q \right) \right) > s^q.$$

Next, notice that there exists $\bar{q}(U_N, \Pi_N)$ such that if $q < \bar{q}(U_N, \Pi_N)$ and $\delta < \bar{\delta}(U_N, \Pi_N)$, then $\pi_{pre}^{q+} \left(u_{pre,E_C,\ell} \left(\bar{u}_{pre,E_C}^q \right) \right) = B/b$ and this inequality is satisfied if

$$p\pi_{pre}^{q+} \left(u_{pre,E_C,h} \left(\bar{u}_{pre,E_C}^q \right) \right) + (1-p)\frac{B}{b} > s^0.$$

The left-hand side of this inequality is weakly bigger than

$$p\pi_{pre}^{q-}(B) + (1-p)\frac{B}{b} \geq p\pi_{post}^-(B) + (1-p)\frac{B}{b},$$

so it suffices to show that $p\pi_{post}^-(B) + (1-p)(B/b) > s^0$.

By construction,

$$s^0 = \frac{b - \pi_{pre}^0(\bar{u}_{E_C})}{B - \bar{u}_{E_C}} \text{ and } \pi_{post}^-(B) = \frac{b - \pi_{post}(\bar{u}_{E_C})}{B - \bar{u}_{E_C}}.$$

Since $\bar{u}_{E_C} = (1-\delta)b + \delta B$, we have that $B - \bar{u}_{E_C} = (1-\delta)(B-b)$. Further,

$$\pi_{post}(\bar{u}_{E_C}) - \pi_{pre}^0(\bar{u}_{E_C}) = (1-p)\delta \left[\pi_{post}(u_{post,E_C,\ell}(\bar{u}_{E_C})) - \pi_{pre}^0(u_{E_C,\ell}(\bar{u}_{E_C})) \right],$$

so the inequality becomes

$$(1-\delta)(B-b)\frac{B}{b} + \pi_{pre}^0(\bar{u}_{E_C}) - b > p\delta \left[\pi_{post}(u_{post,E_C,\ell}(\bar{u}_{E_C})) - \pi_{pre}^0(u_{E_C,\ell}(\bar{u}_{E_C})) \right].$$

By the proof of Proposition 1, $\pi_{pre}^0(\bar{u}_{E_C}) > b$. Finally, we can note that since $U_N > pb$ and $\Pi_N < \left(\frac{1}{p} + p - 1\right)B$,

$$\pi_{post}(u_{post,E_C,\ell}(\bar{u}_{E_C})) - \pi_{pre}^0(u_{E_C,\ell}(\bar{u}_{E_C})) \leq \pi_{post}(pb) - pB < \left(\frac{1}{p} + p - 1\right)B - pB;$$

further, because the incentive constraint holds with equality, and $u_{E_C,h}(\bar{u}_{E_C}) = B$ and $u_{E_C,\ell}(\bar{u}_{E_C}) \leq pb$, we have that

$$\frac{\delta}{1-\delta} \leq \frac{B-b}{B-pb}.$$

Combining these inequalities gives us the desired inequality. ■