The Burden of Past Promises

Jin Li
Northwestern University

Niko Matouschek
Northwestern University

Michael Powell
Northwestern University

February 2014

Preliminary and Incomplete

Abstract

We explore the evolution of a firm's organization and performance. The owner and her employee play an infinitely repeated trust game in which the owner benefits from delegation only if the employee honors her trust by choosing her preferred project. The owner, however, cannot observe whether this project is available. We characterize the optimal relational contract and highlight two implications. First, profits decline over time as the firm's organization evolves from flexibility to rigidity. Second, which type of rigid organization the firm converges to–and thus its long run profitability–is determined by random events in its early history.

1 Introduction

"A good relationship takes time" is popular advice among therapists and advice columnists. A common rationale is that relationships are based on trust and that building trust takes time. Once partners trust each other, they can motivate cooperation by rewarding good behavior today with the promise to take various actions in the future.

At some point, however, the future becomes the present and yesterday's promises become today's obligations. A relationship can then get bogged down by the need to fulfill the very promises that ensured its success early on. Time therefore need not be the friend of a good relationship. Instead, it can be its foe, as many readers and clients of the before-mentioned experts will attest.

In this paper we argue that such dynamics do not only arise in personal relationships but are also a feature of optimally managed relationships between firms and their employees. In particular, we argue that firms often motivate their employees to make "good" decisions by linking their future discretion to their current decisions. Essentially, firms pay their employees with control rather than cash. The optimal allocation of control then needs to balance the desire to influence current decisions with the need to reward or punish past decisions. We show that under the optimal allocation of control, the owner of a firm is able to motivate her employee to make good decisions early on in their relationship. Eventually, however, the owner either has to reward the employee by permanently giving him more discretion than she would if she were not obligated by her past promises, or she has to punish him by giving him permanently less discretion. In either case, the owner is no longer able to make efficient use of the employee's information, the firm's decision making becomes inertial, and its performance declines. We then show that these dynamics speak to both the failure of established firms to adapt to changes in their environments and to the emergence of persistent organizational and performance differences between seemingly similar firms.

The problems that Apple Computer experienced in the 1990s provide an example for the type of dynamics that motivate this paper. As Sull (1999) observes:

"*Managers can also find themselves constrained by their relationships with employees, as the saga of Apple Computer vividly illustrates. Apple's vision of technically elegant computers and its freewheeling corporate culture attracted some of the most creative engineers in the world, who went on to develop a string of smash products including the Apple II, the Macintosh, and the PowerBook. As computers became commodities, Apple knew that its continued health depended on its ability to cut costs and speed up time to market. Imposing the necessary discipline, however, ran counter to the Apple culture, and top management found itself frustrated whenever it tried to exert more control. The engineers simply refused to change their ways. The relationships with creative* *employees that enabled Apple's early growth ultimately hindered it from responding to environmental changes*" (Sull 1999, p.7).

In contrast to the standard intuition that "a good relationship takes time" some relationships therefore get worse over time. And they get worse, in particular, because the discretion firms promise their employees initially ends up constraining their ability to respond to changes in their environments later on.

The model we examine is an infinitely repeated game between the owner of a firm and her employee. At the beginning of the stage game, each party decides whether to enter the relationship. If both do enter, the owner either centralizes–in which case she chooses a status quo project herself–or she delegates–in which case the employee chooses between his preferred project and, if available, the owner's. Each party prefers his or her preferred project to the status quo and the status quo to the other parties' preferred project. The stage game is therefore essentially a trust game. In contrast to a standard trust game, however, only the employee can observe whether the owner's preferred project is available. If the employee chooses his preferred project, the owner therefore does not know if he betrayed her trust or simply had no choice.

Notice that the model rules out monetary transfer. This assumption is the defining characteristic of the literature on delegation that builds on Holmstrom (1977, 1984) (for surveys see, for instance, Bolton and Dewatripont (2013) and Gibbons, Matouschek, and Roberts (2013)). And it is based on the view that a variety of practical factors often make it difficult to pay for decisions.

Even if firms cannot pay for decisions, though, they should be able to motivate decision making through other means. As Prendergast and Stole (1999) observed, for instance: "*A striking characteristic of work life is that one cannot reward individuals in cash for some things, but can compensate them in other ways*" (Prendergast and Stole 1999, p.1007). Similarly, Cyert and March (1963) observed some fifty years ago that "*a significant number of these payments* [within organizations] *are in the form of policy commitments*" (Cyert and March 1963, p. 35). In this paper we argue that these policy commitments often take the form of future control rights. The question then is how the relationship evolves if the owner motivates the employee with the promise of such control rights rather than cash.

To answer this question, we characterize the optimal relational contract, that is, the Perfect Public Equilibrium that maximizes the owner's expected payoff. We show that the owner initially delegates to the employee with the understanding that he chooses her preferred project whenever it is available. To motivate the employee to do so, the owner keeps track of how often he chooses each project. If the employee chooses the owner's preferred project sufficiently often, in a sense that we make precise below, the owner eventually rewards him by delegating to him permanently. And if the employee does not choose her preferred project sufficiently often, the owner instead punishes him by either centralizing permanently, or even terminating the relationship, where the type of punishment if depends on the parameters of the game. In contrast to the well-known equilibria in Green and Porter (1984), therefore, the parties do not alternate between reward and punishment phases. Instead, the owner delays rewards and punishments for as long as possible and then, eventually, administers them with maximum force, that is, permanently.

Inertia and Decline A key feature of the optimal relational contract is that the firm's performance always declines over time. Notice that this is the case even if the employee performs well initially and there is thus no need to punish him. To see this, note that if the employee performs well initially, the owner has to reward him eventually and the optimal way to do so is to delegate to him permanently. Once the owner delegates to him permanently, however, the employee no longer uses his information in the firm's favor and, as a result, its performance declines. Eventually, therefore the relationship between the owner and the employee gets bogged down by the need to fulfill the very promises that ensured its success early on.

The result that the firm inevitably gets worse at making use of the employee's information speaks to the observation that many firms appear to become more inertial over time. Bower and Christensen (1996), for instance, observe 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.*" Similarly, Kreps (1996) argues that "*It is widely held that organizations exhibit substantial inertia in what they do and how they do it (Hannan and Freeman, 1984). In the face of changing external circumstances, organizations adapt poorly or not at all; the economy and/or market evolves as much or more through changes in the population of live organizations than through changes in the organizations that are alive*" (Kreps 1996, p.577). Our model suggests that the inertia of established firms might be the result of the promises that allowed these firms to adapt when they were still young. The flexibility of young firms, and the inertia of established ones, are then two sides of the same coin.

A striking observation that our main model cannot account for is that some firms seem to fail to adapt to information even when that information is publicly available (Schaefer 1998). Sears, for instance, only closed its troubled catalog business after analysts had recommended they do so for many years (Scussel 1991). We explore this issue in our main extension in which we allow for a publicly observable project to become available at a random time. There show that even though the owner would always adopt this project if it were available from the start, she may only do so with some delay, and potentially never do so, if it becomes available later on. The same mechanism that can generate inertia with respect to private information can therefore also generate inertia with respect to public information.

Persistent Organizational and Performance Differences Another key feature of the optimal relational contract is that it generates multiple, long run steady states that are associated with different organizational structures. The model therefore provides a rationale for the widely held view that firms' structures depend on their histories. To once again quote Kreps (1996): "*Organizational policies/procedures tend to be derived from the early history of the organization (Stinchcombe, 1965; Hannan and Freeman, 1977) and to be derived (or at least crystallized out of) specific noteworthy events in the early history of the organization (Schein, 1983)*" (Kreps 1996, p. 577).

Since the different organizational structures are associated with different payoffs, the model also speaks to the observation that there are large and persistent performance differences across firms within narrowly defined industries (for a survey see, for instance, Syverson (2011)). Recent empirical evidence suggests that some of these differences are due to differences in how firms are organized (see, in particular, Bloom et al. (2007, 2013) and, for a survey, Gibbons and Henderson (2012)). If this is so, however, then why don't less successful firms simply imitate the organizational practices of their more successful rivals? After all, such practices are not protected by patents. Our model suggests that one reason may be that firms' histories serve as an endogenous barrier to imitation. One firm may, for instance, be able to centralize decision making without triggering resentment among its employees. In another, and seemingly identical firm, however, employees may view decentralization as their reward for previous achievements.

Transfers and other Loose Ends To conclude the introduction we want to briefly address two assumptions that might otherwise distract the reader. The first is that the owner cannot use any monetary transfers to motivate the employee. This assumption is stronger than what we need for our results. In particular, we show in an extension that as long as the employee is liquidity constrained, the owner could not do any better, and for some discount rates would do strictly worse, if she motivated the employee with cash rather than control.

The second assumption is that we model delegation as a trust game. We do so because it is a well known game that captures the basic problem with delegation. Modeling delegation as a trust game, however, requires an assumption that is not common in the literature on delegation, which is that a particular project–the status quo project–is only available to the owner. One justification for this assumption is that the opportunity to choose the status quo project may be "fleeting" and thus no longer available once the owner has delegated to the employee. Even though we think this is plausible, we also examine an extension in which we allow the status quo project to be available to the employee. We show that permanent centralization is then no longer a long run steady state. The result that the firm's performance always declines over time, and that it does so even if the employee performs well initially, however, continues to hold.

2 The Model

A risk-neutral principal and a risk-neutral agent are in an infinitely repeated relationship. Time is discrete and we denote it by $t = \{1, 2, ...\}$. 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.

Stage Game At the beginning of the stage game, the principal and the agent simultaneously decide whether to enter the relationship. We denote their entry decisions by $e_j \in \{0, 1\}$ for $j = P, A$, where $e_j = 1$ denotes entry. If at least one party decides not to enter, both realize a zero payoff and time moves on to the next period.

If, instead, both parties do decide to enter, the principal next decides whether to delegate the right to choose a project to the agent. We denote the delegation decision by $d \in \{0, 1\}$, where $d = 1$ denotes delegation. Moreover, we denote both projects and project choices by k and the principal's and the agent's stage game payoffs, conditional on both parties having entered the relationship, by $\Pi(k)$ and $U(k)$.

If the principal decides not to delegate to the agent, she chooses a safe project $k = S$ that generates payoffs $\Pi(S) = U(S) = a > 0$. If, instead, the principal does delegate to the agent, the agent can choose between up to two projects. One of these projects is the agent's preferred project $k = A$ and the other is the principal's preferred project $k = P$. The agent's preferred project gives the agent a payoff $U(A) = B$ and the principal a payoff $\Pi(A) = b$, where $B > a > b > 0$. Analogously, the principal's preferred project gives the principal a payoff $\Pi(P) = B$ and the agent a payoff $U(P) = b$. Delegation therefore takes the form of a trust game in which the principal only benefits from delegation if she can trust the agent to choose her preferred project. The assumption that payoffs are symmetric facilitates the exposition but it is not important for our results. We summarize the stage game payoffs in Figure 1.

The key feature of the game is that the principal's preferred project is not always available and that only the agent can observe whether it is available. The principal therefore cannot distinguish a betrayal of her trust from a lack of opportunity to cooperate. In particular, the principal's preferred project is only available with probability $p \in (0,1)$, where the availability is independent across periods. Other than the availability of the principal's preferred project, all information is publicly observable.

Finally, after the parties have realized their payoffs, they observe the realization x of a public randomization device, after which time moves on to the next period.

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

$$\pi_t = (1-\delta)\,\mathrm{E}_t\left[\sum_{\tau=t}^{\infty} \delta^{\tau-t} e_{P,t} e_{A,t} \Pi\left(k_t\right)\right]$$

and the agent's expected payoff is given by

$$u_t = (1-\delta)\,\mathrm{E}_t\left[\sum_{\tau=t}^{\infty} \delta^{\tau-t} e_{P,t} e_{A,t} U\left(k_t\right)\right].$$

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 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. 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 strategy is without loss of generality because our game has only one-sided private information and is therefore a game with the product structure (see, for instance, p.310 in Mailath and Samuelson (2006)). 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 $h_{t+1} = \{e_{P,\tau}, e_{A,\tau}, d_\tau, k_\tau, x_\tau\}_{\tau=1}^{t}$ denote the public history at the beginning of any period $t+1$ and let H_{t+1} denote the set of period $t+1$ public histories. Note that $H_1 = \Phi$. A public strategy for the principal is a sequence of functions $\{E_{P,t}, D_t, K_{P,t}\}_{t=1}^{\infty}$, where $E_{P,t} : H_t \to \{0,1\}$,$D_t : H_t \cup \{e_{P,\tau}, e_{A,\tau}\} \to \{0,1\}$,$K_{P,t} : H_t \cup \{e_{P,\tau}, e_{A,\tau}, d_t\} \to \mathcal{K}_P$, and where $\mathcal{K}_P = \{S\}$ is the set of projects available to the principal. Similarly, a public strategy for the agent is a sequence of functions $\{E_{A,t}, K_{A,t}\}_{t=1}^{\infty}$, where $E_{A,t} : H_t \to \{0,1\}$ and $K_t : H_t \cup \{e_{P,\tau}, e_{A,\tau}, d_t\} \to \mathcal{K}_{A,t}$, and where $\mathcal{K}_{A,t} \in \{\{A\}, \{A,P\}\}$ is the set of projects available to the agent.

We define an "optimal relational contract" as a PPE that maximizes the principal's average per-period payoff. Our goal is to characterize the set of optimal relational contracts.

3 Benchmarks

The model we just described makes three key assumptions: (i.) the stage game is infinitely repeated, (ii.) the principal cannot observe the projects that are available to the agent, and (iii.) transfers are not feasible. We will see below that all three assumptions are crucial for our results. To highlight the role of these assumptions, and to get familiar with the model, we start by considering three benchmarks in which we relax each of the three assumptions in turn.

The Static Game Suppose first that the parties play the stage game only once. The game they play is then essentially a trust game. We say "essentially" because it differs from the standard version of a trust game in two ways. First, before the principal and the agent play the trust game, each has the opportunity to opt out. We allow the parties to opt out since employees can always leave their firms and managers can typically fire their workers. Because the parties can opt out, there is an equilibrium in which neither party enters the relationship. We will see below that, in the repeated game, the parties use this equilibrium to deter publicly observable deviations, such as the principal not delegating to the agent when she is supposed to do so.

The second difference between the stage game and a standard trust game is that the principal cannot observe the actions that are available to the agent. If the game is played only once, this difference is irrelevant since the agent will always betray the principal's trust, no matter what the principal can observe. Anticipating this behavior by the agent, the principal does not trust the agent in the first place. The second equilibrium of the static game is therefore one in which both parties enter the relationship and the principal does not delegate to the agent. This, of course, corresponds to the equilibrium of a standard trust game. And it captures, albeit in a stark way, the view that a principal is more likely to delegate to an agent if she can trust him not to take advantage of his delegated powers.

The Game with Public Information Suppose now that the stage game is infinitely repeated, as in our main model. In contrast to our main model, however, suppose that the principal can observe the projects that are available to the agent. In the Appendix we show that the optimal relational contract then depends on whether the discount factor is above a critical value that lies strictly between zero and one. If the discount factor is below the critical value, the principal cannot do better than to centralize in every period. If it is above the critical value, however, the principal can do better by having both parties play standard trigger strategies. Under these strategies, the principal starts out by delegating to the agent with the understanding that he will choose the principal's preferred project whenever it is available. The principal will continue to do so unless the agent ever violates this understanding, in which case she opts out of the relationship in all future periods. In response, the agent chooses the principal's preferred project whenever it is available. In the game with public information, there therefore always exists an optimal relational contract that is stationary and does not involve any punishment on the equilibrium path. We will see below that this is not the case in our main model, in which the principal cannot observe the projects that are available to the agent.

The Game with Transfers Suppose now that the principal cannot observe the projects that are available to the agent, as in our main model. In contrast to our main model, however, suppose that the principal can use monetary transfers to motivate the agent. In particular, suppose that at the beginning of the stage game, the principal can make a take-it-or-leave-it offer to the agent in which she can contractually commit to a fixed wage and promise to pay a bonus. In the Appendix we show that, as in the game with public information, the optimal relational contract depends on whether the discount rate is above a critical value that lies strictly between zero and one. If the discount rate lies below the critical value, the principal cannot do better than to centralize in every period. If it lies above the critical value, however, the principal can do better by having both parties play standard trigger strategies. The principal again starts out by delegating to the agent. In contrast to the game with public information, however, she now offers to "pay" him a wage equal to $-B$ and promises to pay him a bonus equal to $(B-b)$ whenever he chooses the principal's preferred project. In response, the agent accepts the offer and chooses the principal's preferred project whenever it is available, unless the principal ever reneges on her promise to pay the bonus, in which case the agent opts out of the relationship in every future period. In the game with transfers, as in the game with public information, there therefore always exists an optimal relational contract that is stationary and does not involve any punishment on the equilibrium path. As mentioned above, this is not the case in our main model, in which the principal cannot rely on transfers to motivate the agent.

4 Preliminaries

In this section, we characterize the PPE payoff set. We first list the constraints that payoffs have to satisfy to be in the PPE payoff set. In Section 4.2 we then derive a constrained maximization problem that characterizes the payoff frontier and show that it fully determines the optimal relational contract. In Section 5 we can then characterize the optimal relational contract by solving this problem.

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 between two equilibrium payoff pairs that are each generated by pure actions. There are four sets of pure actions. First, both parties enter the relationship after which the principal delegates to the agent with the understanding that he chooses the principal's preferred project whenever it is available. We call this set of actions "cooperative delegation" and denote it by D_C. Second, both parties enter the relationship after which the principal delegates to the agent with the understanding that he can always choose his preferred project. We call this action "uncooperative delegation" and denote it by D_U. Third, both parties enter the relationship after which the principal centralizes and chooses the safe project. We call this action "centralization" and denote it by C. Finally, neither party enters the relationship. We call this set of actions "exit" and denote it by E. 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 four sets of pure actions. We then conclude the section by stating the constraint that has to be satisfied if the payoff pair is generated by randomization.

Centralization 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 that the parties realize under centralization therefore have to satisfy the self-enforcement constraint

$$(u_C, \pi_C) \in \mathcal{E}. \tag{SE$_\mathrm{C}$}$$

(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. There is no loss of generality in assuming that if an off-schedule deviation occurs, the parties terminate the relationship by opting out in all future periods, as this is the worst possible equilibrium that gives each party its minmax payoff.

The principal and the agent can deviate off-schedule by opting out of the relationship. If either party does so, he or she realizes a zero payoff this period and in all future periods. Since the payoffs from the three projects are strictly positive, the parties therefore do not have an incentive to deviate off-schedule by opting out of the relationship.

The principal could also deviate off-schedule by delegating. There is no loss of generality in assuming that the agent will then choose his preferred project. By deviating, the principal would therefore reduce her current payoff from a to $b < a$, after which she would make a zero payoff in all future periods. The principal therefore never wants to deviate off-schedule by delegating.

On-schedule deviations are deviations that 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 in the case of 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

$$\pi = (1-\delta)\, a + \delta \pi_C \tag{$\mathrm{PK_C^P}$}$$

and

$$u = (1-\delta)\, a + \delta u_C \tag{$\mathrm{PK_C^A}$}$$

ensure that this is the case.

Cooperative Delegation A payoff pair (u, π) can be supported by cooperative delegation if the following constraints are satisfied.

(i.) Feasibility: For the continuation payoffs to be feasible, they also need to be PPE payoffs. Let (u_ℓ, π_ℓ) denote the parties' continuation payoffs if the agent chooses his preferred project and let (u_h, π_h) denote their payoffs if he chooses the principal's preferred project. The self-enforcement constraint is then given by

$$(u_h, \pi_h)\,, (u_\ell, \pi_\ell) \in \mathcal{E}, \tag{$\mathrm{SE_{D_C}}$}$$

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

(ii.) No Deviation: As in the case of centralization, the principal and the agent never want to deviate off-schedule by opting out of the relationship since doing so gives them a zero payoff. The principal may, however, want to deviate off-schedule by not delegating to the agent, in which case she realizes payoff a this period and a zero payoff in all future periods. To ensure that she does not want to do so, the reneging constraint

$$p\left[(1-\delta)\, B + \delta \pi_h\right] + (1-p)\left[(1-\delta)\, b + \delta \pi_l\right] \geq (1-\delta)\, a \tag{$\mathrm{NR_{D_C}}$}$$

has to be satisfied.

Since the principal does not have any private information, she cannot engage in any on-schedule deviations. The agent, however, may choose his preferred project when the principal's preferred project is available. To ensure that he does not want to do so, the incentive constraint

$$(1-\delta)\,b+\delta u_h \geq (1-\delta)\,B+\delta u_\ell \qquad (\mathrm{IC}_{\mathrm{D_C}})$$

has to be satisfied.

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

$$\pi = p\,[(1-\delta)\,B+\delta\pi_h]+(1-p)\,[(1-\delta)\,b+\delta\pi_\ell] \qquad (\mathrm{PK}^{\mathrm{P}}_{\mathrm{D_C}})$$

and

$$u = p\,[(1-\delta)\,b+\delta u_h]+(1-p)\,[(1-\delta)\,B+\delta u_\ell]\,. \qquad (\mathrm{PK}^{\mathrm{A}}_{\mathrm{D_C}})$$

Uncooperative Delegation A payoff pair (u,π) can be supported by uncooperative delegation if the following constraints are satisfied.

(i.) Feasibility: We denote the continuation payoffs under uncooperative delegation by (u_{D_U},π_{D_U}). The self-enforcement constraint is then given by

$$(u_{D_U},\pi_{D_U})\in\mathcal{E}. \qquad (\mathrm{SE}_{\mathrm{D_U}})$$

(ii.) No Deviation: As in the case of cooperative delegation, the principal and the agent never want to deviate off-schedule by opting out of the relationship since doing so gives them a zero payoff both this period and in all future periods. The principal may, however, want to deviate off-schedule by not delegating to the agent. To ensure that she does not want to do so, the reneging constraint

$$(1-\delta)\,b+\delta\pi_{D_U} \geq (1-\delta)\,a. \qquad (\mathrm{NR}_{\mathrm{D_U}})$$

has to be satisfied.

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

$$\pi = (1-\delta)\,b+\delta\pi_{D_U}. \qquad (\mathrm{PK}^{\mathrm{P}}_{\mathrm{D_U}})$$

for the principal and

$$u = (1-\delta)\,B+\delta u_{D_U} \qquad (\mathrm{PK}^{\mathrm{A}}_{\mathrm{D_U}})$$

for the agent.

Exit A payoff pair (u, π) can be supported by exit if the following constraints are satisfied.

(i.) Feasibility: We denote the continuation payoffs under centralization by (u_E, π_E). The self-enforcement constraint is then given by

$$(u_E, \pi_E) \in \mathcal{E}. \tag{SE$_\mathrm{E}$}$$

(ii.) No Deviation: The principal and the agent can deviate off-schedule by entering the relationship. If the principal or the agent does so, he or she realizes a zero payoff this period and in all future periods. The parties therefore do not have an incentive to deviate by entering the relationship. There are no other off- or on-schedule deviations in this case

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

$$\pi = \delta \pi_E \tag{PK$_\mathrm{E}^\mathrm{P}$}$$

for the principal and

$$u = \delta u_E. \tag{PK$_\mathrm{E}^\mathrm{A}$}$$

for the agent.

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

$$(u, \pi) = \alpha (u_1, \pi_1) + (1 - \alpha)(u_2, \pi_2)$$

for some $\alpha \in (0, 1)$.

4.2 The Constrained Maximization Problem

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

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

where $\mathcal{E}$ is the PPE payoff set. We also define

$$\underline{u} = \inf\{u : (u, \pi) \in \mathcal{E}\}$$

and

$$\overline{u} = \sup\{u : (u, \pi) \in \mathcal{E}\}$$

as the smallest and the largest PPE payoff for the agent.

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

LEMMA 1. *The PPE payoff set $\mathcal{E}$ has the following properties: (i.) it is compact, (ii.) $\pi(u)$ is concave, (iii.) the payoff pair (u, π) belongs to $\mathcal{E}$ if and only if $u \in [0, B]$ and $\pi \in [bu/B, \pi(u)]$.*

The first part of the lemma shows that the PPE payoff set is compact. This result follows immediately from the assumption that there is only a finite number of actions. And it implies that for any $u \in [\underline{u}, \overline{u}]$ 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 smallest PPE payoff for the agent is zero and the largest is B. It also shows that, for any $u \in [0, B]$, the smallest PPE payoff for the principal is bu/B and that, for any $\pi \in [bu/B, \pi(u)]$, the payoff pair (u, π) is in the PPE payoff set.

A key implication of the first lemma is that to describe the PPE payoff set, we only need to characterize its frontier. To do so, we need to determine, for each $(u, \pi(u)) \in \mathcal{E}$, whether it is supported a pure action $j \in \{C, D_C, D_U, E\}$ 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 satisfy

$$\delta u_C(u) = u - (1 - \delta) a.$$

(ii.) If the payoff pair is supported by cooperative delegation, the agent's continuation payoff can be chosen to satisfy

$$\delta u_\ell (u) = u - (1-\delta) B$$

and

$$\delta u_h (u) = u - (1-\delta) b.$$

(iii.) If the payoff pair is supported by uncooperative delegation, the agent's continuation payoff satisfy

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

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

$$\delta u_E (u) = u.$$

In the cases of centralization, uncooperative delegation, and exit, the agent's continuation payoffs follow directly from the promise-keeping constraints $\mathrm{PK}^{\mathrm{A}}_{\mathrm{C}}$ and $\mathrm{PK}^{\mathrm{A}}_{\mathrm{D_U}}$. In the case of cooperative delegation, instead, the agent's continuation payoffs follow directly from combining the promise-keeping constraints with the agent's incentive constraint $\mathrm{IC}_{\mathrm{D_C}}$, where we take the incentive constraint to be binding. To see that we can do so, suppose that the incentive constraint is not binding. We can then reduce u_h and increase u_ℓ in such a way that u remains the same, and all the relevant constraints continue to be satisfied. Since the PPE frontier is concave, such a change makes the principal weakly better off.

Next we can use Lemmas 2 and 3 to derive explicit 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 $j \in \{C, D_C, D_U, E\}$ be the highest payoff to the principal given action j and agent's payoff u. We then have

$$\pi_C (u) = (1-\delta) a + \delta\pi (u_C (u)),$$

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

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

and

$$\pi_E (u) = \delta\pi (u_E (u)).$$

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 $u \in [0, B]$*

$$\pi(u) = \max_{q_j \geq 0, u_j \in [0,B]} \sum_{j \in \{C, D_C, D_U, E\}} q_j \pi_j(u_j)$$

such that

$$\sum_{j \in \{C, D_C, D_U, E\}} q_j = 1$$

and

$$\sum_{j \in \{C, D_C, D_U, E\}} q_j u_j = u.$$

The lemma shows that any payoff pair on the frontier is generated either by a pure action j–in which case the weight q_j is equal to one–or by randomization–in which case q_j is less than one. We obtain the frontier by choosing the weights optimally. Notice that the frontier can be thought of as a fixed point to an operator. We show in the proof that the fixed point is unique even though the operator is not a contraction mapping. In the next section, we solve the problem in the lemma to characterize the PPE frontier and thus the optimal relational contract.

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 $\underline{u}_{CD} \in (a, \delta a + (1-\delta) B)$ and $\bar{u}_{CD} = (1-\delta) b + \delta B$ such that the PPE payoff frontier $\pi(u)$ is divided into four regions:*

(i.) For $u \in [0, a]$, $\pi(u) = u$, and $(u, \pi(u))$ is supported by randomization between exit and centralization.

(ii.) For $u \in [a, \underline{u}_{CD}]$, $\pi(u) = ((\underline{u}_{CD} - u) a + (u - a) \pi(\underline{u}_{CD})) / (\underline{u}_{CD} - a)$ and $(u, \pi(u))$ is supported by randomization between centralization and cooperative delegation.

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

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

We illustrate the lemma in Figure 1. 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 randomizations. 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.

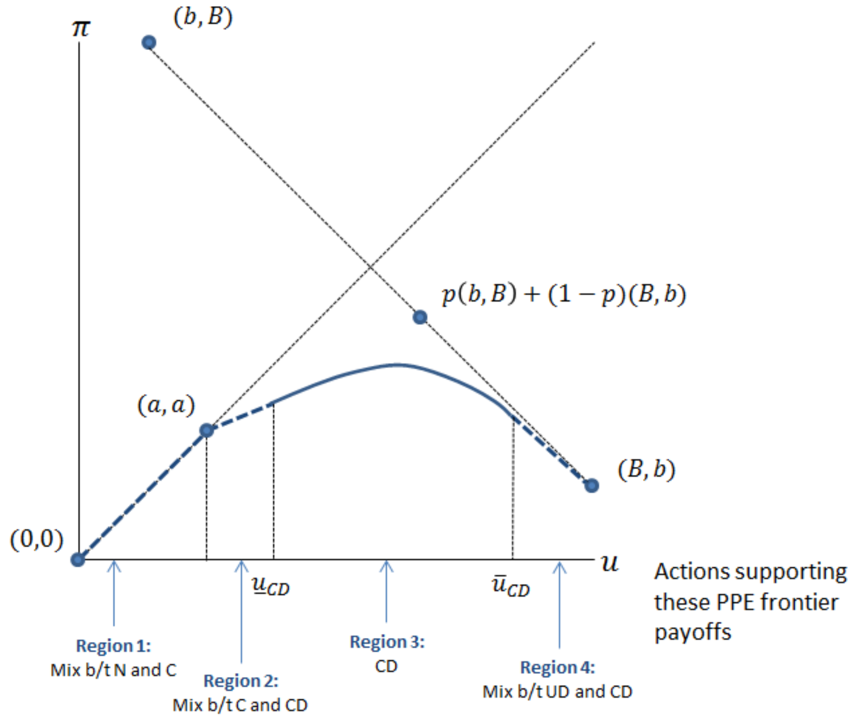

Figure Y: This figure illustrates the feasible stage-game payoffs, the PPE payoff frontier, and the actions that support each point on the frontier. The dotted linear segments are supported by public randomization between their two endpoints, and this public randomization occurs at the end of the period.

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

PROPOSITION 1. *First period: The agent's and the principal's payoffs are given by* $u^* \in [\underline{u}_{CD}, \bar{u}_{CD}]$ *and* $\pi(u^*) = \pi_{D_c}(u^*)$. *The parties engage in cooperative delegation. If the agent chooses the principal's preferred project, his continuation payoff is given by*

$$u_h(u^*) = (u^* - (1-\delta)b)/\delta > u^*.$$

If, instead, the agent chooses his own preferred project, his continuation payoff is given by

$$u_\ell(u^*) = (u^* - (1-\delta)B)/\delta < u^*.$$

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

(i.) If $u = 0$, the parties exit. The agent's continuation payoff is given by $u_E(0) = 0$.

(ii.) If $u = a$, the parties engage in centralization. The agent's continuation payoff is given by $u_C(a) = a$.

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

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

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

To see how the principal optimally increases the agent's continuation value, suppose the agent chooses the principal's preferred project for a number of consecutive periods. The principal then continues to engage in cooperative delegation, and the agent's continuation value continues to increase, until the parties reach a period in which the continuation value passes the threshold $\bar{u}_{CD}$. 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 delegation or she moves to uncooperative delegation, that is, she allows the agent to choose his preferred project even if her preferred project is available. Finally, once the principal has moved to uncooperative delegation, she stays there in all subsequent periods.

To see how the principal optimally decreases the agent's continuation value, suppose instead that the agent chooses his own preferred project for a number of consecutive periods. The principal then continues to engage in cooperative delegation, and the agent's continuation value continues to decrease, until the parties reach a period in which the continuation value falls below the threshold $\underline{u}_{CD}$. 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 $u \in [a, \underline{u}_{CD}]$, the principal either continues to engage in cooperative delegation or she moves to centralization. And if, instead, $u \in [0, a)$, the principal either moves to centralization or she exits the relationship in the next period. Finally, once the principal has moved to either centralization or exit, she stays there in all subsequent periods.

A key feature of the optimal relational contract is that once the principal chooses an action other than cooperative delegation, she takes that action in all future periods. It is therefore not optimal for the parties to cycle between reward and punishment phases, as in the well known class of equilibria that Green and Porter (1984) first introduced. To see why such equilibria are not optimal, notice that both rewards–letting the agent choose his preferred project even when the principal's is available–and punishments–opting out or centralizing–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. Because of this benefit, the principal wants to delay them as much as she can.

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 termination or centralization? The next proposition answers these questions.

PROPOSITION 2: *In the optimal relational contract, the principal chooses cooperative delegation for the first τ periods, where τ is random and finite with probability one. Moreover, there exists a threshold p^* such that the relationship never terminates if $p \leq p^*$. If, instead, $p > p^*$, punishment can take the form of either termination or centralization, depending on the history of the relationship.*

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 termination or centralization–is that it depends on the probability p that the principal's preferred project is available. Having characterized the optimal relational contract, we now turn to its implications, which we already sketched and discussed in the introduction.

The first implication is that the principal's payoff declines over time, even if the relationship does not terminate. 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 permanent centralization–in which case the principal makes $a < \pi(u^*)$–or permanent delegation–in which case she makes $b < \pi(u^*)$. The principal's payoff declines over time, because the firm gets worse at using the agent's information. And the firm gets worse at using the agent's information because, eventually, the principal either has to reward the agent–by letting him choose any project–or punish him–by choosing a project herself. In either case, the firm's decision no longer reflect the agent's information.

The second implication is that the organizational structure that the firm converges to, and thus the long run payoff that the principal realizes, depend on random events in the firm's early history. In particular, whether the firm converges to permanent centralization–in which case the principal's payoff is given by a–or whether it converges to permanent decentralization–in which case it is given by $b < a$–depends on the randomly determined availability of projects in the periods before the firm converges to either organization. This suggests that persistent performance differences across seemingly identical firms may be due to persistent organizational differences which, in turn, may be due to random differences in the early history of those firms.

Also, and related, the model suggests an explanation for why some under-performing firms do not copy the organizational practices of their more successful rivals, even though such practices are not protected by patents. In particular, it suggests that such firms may not imitate their more successful rivals since their seemingly inefficient organizations are either a reward for past successes or a punishment for past failures. In either case, employees would view the adoption of a different organizational structure as the violation of a mutual understanding and punish the firm accordingly. A firm's history can therefore serve as a barrier to organizational imitation.

6 Failure to Adapt to Public Information

Business history is littered with established firms that failed to adapt to changes in their environments (see, for instance, Bower and Christensen (1996)). In the previous section we argued that this failure to adapt may be the price that established firms have to pay for their ability to adapt when they were still young. A key feature of our main model, however, is that the changes in the environment that the firm fails to adapt to are privately observed by the employee. As such, our main model cannot account for the apparent failure of some firms to adapt to changes even when they are publicly observable. In the Introduction we mentioned Sears' failure to close its catalog business as one example (Scussel 1991). Schaefer (1998) discusses another example:

"*As an example* [of firms that are unable to make seemingly obvious changes until the survival of the organization is threatened] *consider the differences in product development processes between General Motors and Toyota. For years prior to 1991, auto-industry analysts had highlighted Toyota's insistence on "design for manufacturability" as an important cost-saving device. GM, although burdened with the highest costs and slowest product development processes in the industry, retained an organizational structure in which the design group operated autonomously. Only after record-breaking losses of $4.5 billion in 1991 and $23.5 billion in 1992 did GM undertake a series of changes that forced designers to report directly to engineering.*" (Schaefer 1998, p. 251-252).

In line with our main model, this discussion suggests that GM failed to adopt more efficient design practices since it had delegated those choices to the design group. In contrast to our main model, the availability of those more efficient practices was not privately observed by that group and, instead, had been publicized by analysts for years.

In this section, we show that the same forces that cause the firm to become inertial with respect to the employee's private information can also make it inertial with respect to information that is publicly available. In particular, we show that if the owner pays the employee with control rather than cash, she may find it optimal to promise that she will delay the adoption of a profitable project that will become available at some point in the future. And she may do so, even when the owner can observe when the project has become available and, once it has, could simply choose it herself.

Specifically, we now consider an extension of our main model in which the periods are divided 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 the main model is that the owner can now choose between two projects: the status quo and a new project that gives the owner a payoff π_N and the employee a payoff u_N. And the only difference between the stage game in the pre-opportunity phase and the one in the main section 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 $q \in (0,1)$, which is independent across periods.

To make the analysis interesting, we assume that the new project is neither too attractive nor too unattractive to both parties. Specifically, we assume that the employee's payoff from the new project u_N satisfies $u_N \in (\underline{u}_{CD}, \bar{u}_{CD})$, where $\underline{u}_{CD}$ and $\bar{u}_{CD}$ are the threshold values of u between which cooperative delegation is optimal in our main model. This assumption ensures that there is a trade-off between motivating the employee and choosing the new project. We also assume that $\pi_N \in (\pi^*, \pi(u_N) + D]$ for some $D > 0$, where π^* is the principal's highest equilibrium payoff in the main model. The assumption that $\pi_N > \pi^*$ ensures that the new project is sufficiently profitable so that, if it were available in the first period, the owner would choose it immediately. The assumption that $\pi_N \leq \pi(u_N) + D$, in turn, ensures that the project is not so profitable that the owner would always choose it, no matter when it becomes available.

Notice that the game is now a random game rather than a repeated one. To characterize the optimal relational contract, we therefore have to characterize two payoff frontiers: $\pi_{\text{Pre}}(\cdot)$–the payoff frontier in the pre-opportunity phase–and $\pi_{\text{Post}}(\cdot)$–the frontier in the post-opportunity phase. Since the game transitions from the pre- to the post-opportunity phase, but not the reverse, we first characterize $\pi_{\text{Post}}(\cdot)$ and then $\pi_{\text{Pre}}(\cdot)$.

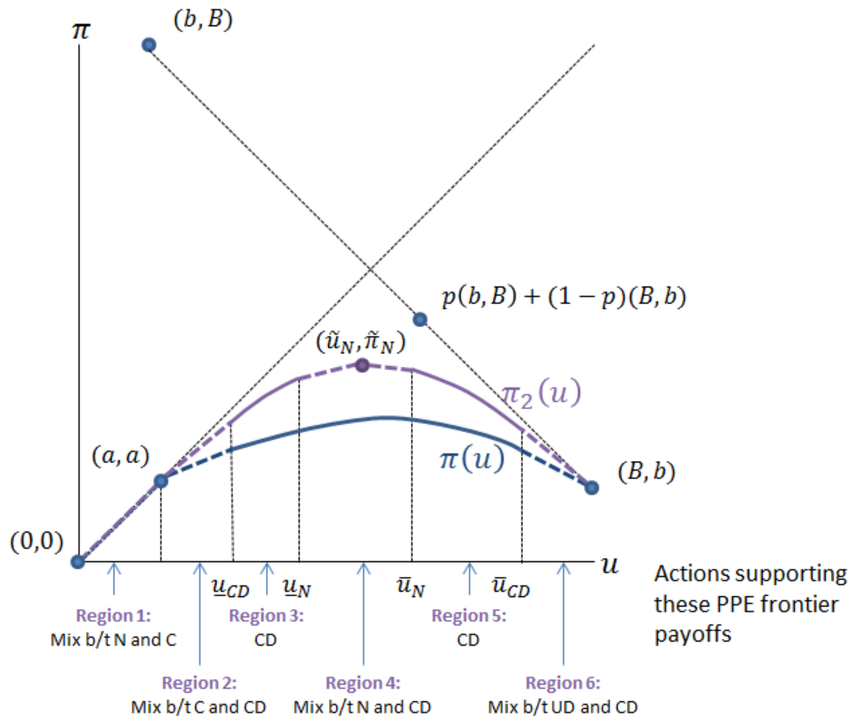

Figure X: This figure illustrates the PPE payoff frontier for the baseline model and the PPE payoff frontier as well as the actions that support each point on the frontier for phase 2 of the current extension.

We characterize $\pi_{\text{Post}}(\cdot)$ formally in Appendix B. Figure x illustrates its main features and compares them to those of the payoff frontier $\pi(\cdot)$ in our main model. One difference is that $\pi_{\text{Post}}(\cdot)$ is everywhere above $\pi(\cdot)$. This reflects the fact that the owner's payoff from the new project π_N is higher than her highest equilibrium payoff in the main model. In fact, π_N is highest payoff on the post-opportunity frontier. As we claimed above, the owner would therefore always choose the new project if it were available in the first period. Another difference between the payoff frontiers is that the payoffs from the new project (u_N, π_N) are on $\pi_{\text{Post}}(\cdot)$ but not $\pi(\cdot)$. This is the case since choosing the new project can be sustained as a sub-game perfect equilibrium of the stage game in the post-opportunity phase but obviously not in the main model. Finally, for values of u around u_N, the frontier payoffs in the post-opportunity phase are not supported by cooperative delegation as they are in the main model. Instead, they are supported by randomization between cooperative delegation and choosing the new project.

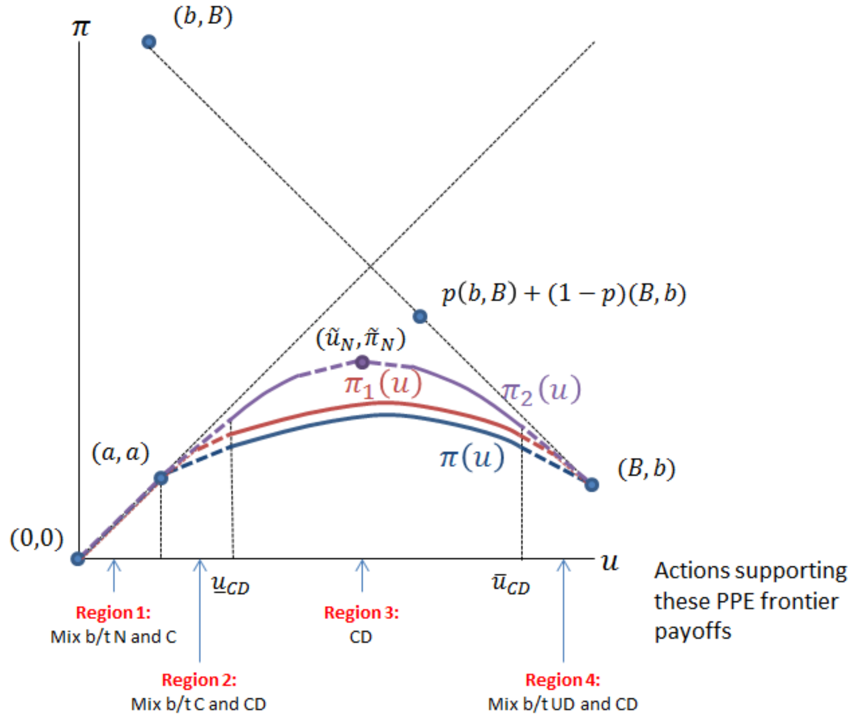

Figure W: This figure illustrates the PPE payoff frontiers for the baseline model and for phases 1 and 2 of the current extension. It also describes the actions that support each point on the phase-1 frontier.

Consider now the payoff frontier in the pre-opportunity phase. We illustrate its main features in Figure Y and once again relegate the formal analysis to Appendix B. The figure shows the payoff frontier in the pre-opportunity phase looks very similar to the one in the main model. Recall, however, that in the main model, payoffs on the frontier are supported by continuation payoffs that are again on the same frontier. In the pre-opportunity phase, in constrast, 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 employee receives if the new project becomes available away from u_N. From Figure x, however, it then follows that there is at least some chance that the owner will not choose the new project as soon as it becomes available and may, in fact, never do so. Our next proposition provides conditions under which this is indeed the case.

PROPOSITION 3: *For each* u_N,

(i.) *There exists a* $\overline{\pi}(u_N)$ *such that for all* $\pi_N \in (\pi(u_N), \overline{\pi}(u_N))$, *there exists a history* h^T *such* *that* $\Pr\left(u_T = \tilde{u}_N | h^T\right) < 1$*, where* T *is the first period in the post-opportunity phase and* h^T *is the history of public outcomes up to period* T.

(ii.) *There exists a* $\widehat{\pi}(u_N) \leq \overline{\pi}(u_N)$ *such that for all* $\pi_N \in (\pi(u_N), \widehat{\pi}(u_N))$*, there exists history of path* h^T *such that* $\Pr\left(u_t = u_N | h^T\right) = 0$ *for all* $t \geq T$.

The first part of the proposition provides conditions under which the owner does not choose the new project as soon as it becomes available. Suppose, for instance, that the employee has chosen the owner's preferred project sufficiently often in the pre-opportunity phase so that his continuation payoffs (in case the new project becomes available) exceed u_N. The owner is then rewarding the employee 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 owner may in fact promise never to adopt the new project. Suppose, for instance, that the employee has chosen the owner's preferred project so often during the pre-opportunity phase that his continuation payoff (in case the new preoject beomces available) does not only exceed u_N but is actually equal to B. The owner is then rewarding the employee for his excellent performance in the pre-opportunity phase by promising him that he will always be able to choose his own preferred project, even if the new project becomes available.

7 Extensions

In this section, we examine two of the assumptions that our main results are based on. The first is that no transfers are allowed between the principal and the agent. The second is that the safe project is fleeting: when the principal delegates to the agent, the agent cannot choose the safe project.

7.1 Transfers from Principal to Agent

Suppose that at the end of period t, the principal can pay the agent a non-negative transfer $w_t \geq 0$ contingent upon the agent's project choice. The relational contract therefore specifies a bonus scheme and an action to be taken in each period. Denote by $\pi_T(u)$ the PPE payoff frontier of this extended game with transfers. The main result in this section is that allowing transfers from the principal to the agent does not affect the results of Propositions 2 or 3.

PROPOSITION 4: $\pi_T(u) = \pi(u)$. Moreover, the optimal relational contract specified in the game with no transfers is also an optimal relational contract when transfers from the principal to the agent are allowed.

In the baseline model, the agent is rewarded for choosing the principal's preferred project through an increase in the probability that he will be able to choose his own project indefinitely in the future. When the agent chooses his own preferred project, he is punished through an increase in the probability that the safe action will be chosen indefinitely in the future. Since both cooperative delegation and uncooperative delegation yield a total surplus of $b + B$ in each period, rewards are simply a reallocation of total surplus from the principal to the agent, while punishments actually result in a decrease in total surplus.

If unrestricted transfers from the agent to the principal were allowed, then surplus-destroying punishments could be avoided by requiring such transfers when the agent chooses his own preferred project. Consequently, the principal would be indifferent between rewarding the agent through increases in his continuation payoff or through monetary transfers. However, when unrestricted transfers from the agent to the principal are not possible, punishment requires surplus destruction, and it requires more surplus destruction the lower is the agent's continuation payoff, because the payoff frontier is concave. Conversely punishments are less costly the greater is the agent's continuation payoff, and it follows that rewarding the agent with an increase in his continuation payoff is preferable to rewarding him with money.

This result is obtained in part because total surplus under the agent's preferred action is the same as under the principal's preferred action. If, instead, payoffs for the agent's preferred action are (B_A, b_A) and for the principal's preferred action are (b_B, B_B) with total surplus lower under the agent's preferred action (i.e., $B_A + b_A < B_P + b_P$), then rewarding the agent with an increase in his continuation payoff results in surplus destruction. If transfers are costly, so that it costs the principal $1 + \kappa$ dollars to transfer 1 dollar to the agent, then as long as $1 + \kappa \geq \left| \frac{B_P - b_A}{B_A - b_P} \right|$, monetary transfers will not be used in an optimal relational contract.

7.2 Safe Project Always Available

In the baseline model, whenever the principal delegates to the agent, the agent does not have the option of choosing the safe project. That is, the set of projects that the agent can choose from in period t is $\mathcal{K}_{A,t} \in \{\{A\}, \{A, P\}\}$. As a result, punishment takes the form of centralization: the principal does not delegate, and the principal chooses the safe project. But if the agent has the option to choose the safe project, then punishment may instead take the form of **constrained delegation** in which the principal delegates to the agent with the understanding that if the principal's preferred project is not available, the agent will choose the safe project rather than his own preferred project.

Formally, suppose that $\mathcal{K}_{A,t} \in \{\{S, A\}, \{S, A, P\}\}$, where $\Pr[\mathcal{K}_{A,t} = \{S, A, P\}] = p$. There are now five sets of actions that may be used in equilibrium. In addition to "cooperative delegation," "uncooperative delegation," "centralization," and "exit," the players may also choose "constrained delegation," in which the principal delegates to the agent with the understanding that the agent will choose the principal's preferred project whenever it is available and will choose the safe project otherwise. We denote constrained delegation by $j = CDD$. When parties engage in constrained delegation, and the agent's promised utility is u, the promise-keeping constraint is

$$u = (1-p)\left((1-\delta)a + \delta u_{CDD,\ell}(u)\right) + p\left((1-\delta)b + \delta u_{CDD,h}(u)\right),$$

and the agent will choose the principal's preferred project whenever it is available as long as

$$(1-\delta)b + \delta u_{CDD,h}(u) \geq (1-\delta)a + \delta u_{CDD,\ell}(u).$$

Consequently, continuation payoffs can be chosen to make this inequality hold with equality:

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

We denote the PPE payoff frontier of this extended game by $\pi_S(u)$. Lemma 7 in the appendix characterizes the frontier of the game when the safe project is always available, and figure W below illustrates a number of its properties.

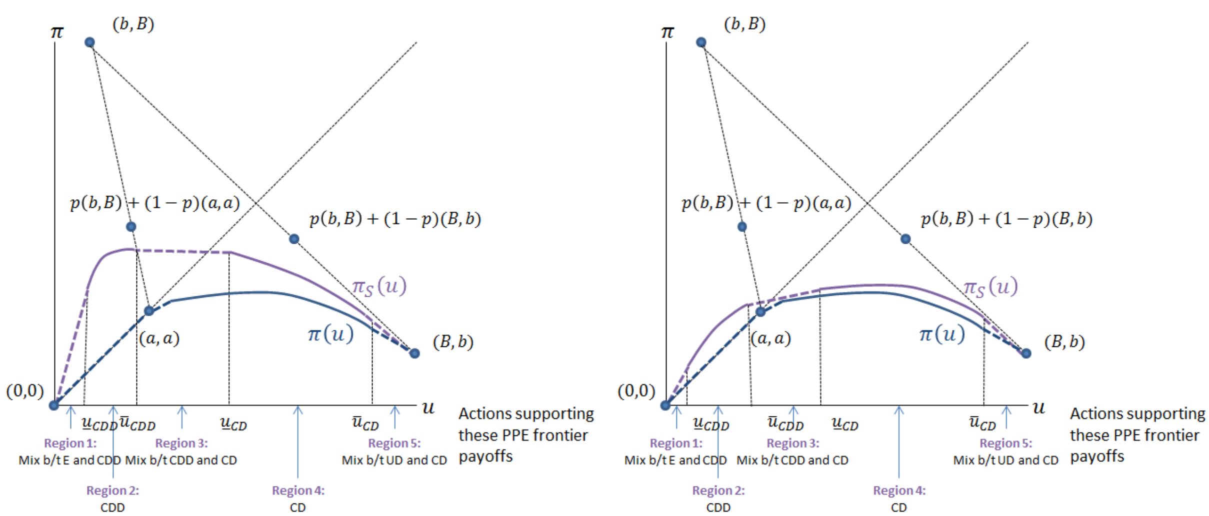

Figure W: These figures illustrate the PPE payoff frontiers for the baseline model and for the extended model in which the safe project can be chosen by the agent. They also describe the actions that support each point on the frontier. The figure on the left illustrates the PPE frontier and actions for the case in which constrained delegation is chosen in the first period. The figure on the right illustrates the PPE frontier and actions for the case in which cooperative delegation is chosen in the first period.

Delegating to the agent with the understanding that the agent will choose the safe project whenever the principal's preferred project is not available yields higher profits for the principal than centralization does. As a result, the PPE payoff frontier lies above the point (a, a), and centralization will never be chosen in equilibrium. The frontier consists of five regions. In the first region, the parties randomize between exit and constrained delegation in the next period. In the second region, parties engage in constrained delegation. In the third region, parties randomize between constrained delegation and cooperative delegation in the next period. In the fourth region, parties engage in cooperative delegation, and in the fifth region, they randomize between cooperative delegation and uncooperative delegation. It may be the case that the highest point on the frontier lies in region 2 or in region 4, depending on the parameters of the model.

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

PROPOSITION 5. *When delegative control is allowed, the optimal relational contract satisfies the following.*

First period: *The agent's and the principal's payoffs are given by* $u^* \in [\underline{u}_{CDD}, \bar{u}_{CDD}] \cup [\underline{u}_{CD}, \bar{u}_{CD}]$ *and* $\pi(u^*) = \max\{\pi_{D_{CD}}(u), \pi_{D_C}(u^*)\}$. *The parties engage in either delegative control or cooperative delegation. In either case, if the agent chooses the principal's preferred project, his continuation payoff increases and drops otherwise.*

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

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

(ii.) If $u \in [\underline{u}_{CDD}, \bar{u}_{CDD}]$*, the parties engage in constrained delegation. If the agent chooses the principal's preferred project, his continuation payoff is given by* $u_h^{D_{CD}}(u) > u$. *If, instead, control is used, his continuation payoff is given by* $u_\ell^{D_{CD}}(u) < u$.

(iii.) If $u \in [\underline{u}_{CD}, \bar{u}_{CD}]$*, the parties engage in cooperative delegation. If the agent chooses the principal's preferred project, his continuation payoff is given by* $u_h^{D_C}(u) > u$. *If, instead, the agent chooses his own preferred project, his continuation payoff is given by* $u_\ell^{D_C}(u) < u$.

(iv.) If $u = B$*, the parties engage in uncooperative delegation. The agent's continuation payoff is given by* $u_{D_U}(B) = B$.

The proposition shows that the principal starts out by engaging in either constrained delegation or cooperative delegation. As in the baseline model, to motivate the agent to choose her preferred project whenever it is available, the principal increases his continuation value whenever he chooses her preferred project, and she decreases his continuation value when he chooses his own preferred project (under cooperative delegation) or the safe project (under constrained delegation). The short-run dynamics are similar whether the optimal relational contract begins with constrained delegation or with cooperative delegation.

When the agent is able to choose the safe project, the principal has two tools available to punish the agent. Under cooperative delegation, if the agent has chosen his own preferred project sufficiently often, his continuation payoff falls, and eventually the principal has to alter his choice of project in order to reduce the agent's per-period payoff. She does so by effectively restricting the set of projects that the agent can choose from and removing the agent's preferred project from this set. Reduced discretion can therefore be a punishment for poor performance under cooperative delegation.

Similarly, under constrained delegation, if the agent has chosen the safe project sufficiently often, his continuation payoff falls. Eventually, the principal has to punish the agent, and she does so by choosing exit with some positive probability. If the agent has chosen the principal's preferred project sufficiently often, his continuation payoff increases, and the principal eventually rewards him with increased discretion, allowing him to choose his own preferred project when her preferred project is not available. Finally, as in the baseline model, the possibility of the relationship moving into uncooperative delegation serves as a potential reward for the agent. We now describe the long-run dynamics in the following proposition.

PROPOSITION 6: *In the optimal relational contract, the principal chooses either cooperative delegation or constrained delegation for the first τ periods, where τ is random and finite with probability one. For $t > \tau$, the relationship results in either termination or entrenchment, depending on the history of the relationship. Both possibilities occur with positive probability for all $p \in (0, 1)$.*

Proposition 6 shows that, as in the baseline model, the relationship eventually settles into one of two steady states: termination or entrenchment. Since centralization is never chosen in equlibrium, the relationship can never settle into permanent centralization.

8 Conclusions

Firms often motivate decision making by their employees with control rather than cash. In choosing the allocation of control, such firms then have to balance the desire to influence current decision making with the need to reward or punish past decision making. In this paper we explored the allocation of control that strikes this balance optimally. For this purpose, we developed a simple, dynamic model of delegation and showed that under the optimal allocation of control, the owner of a firm is able to motivate her employees to make good decisions early on in their relationship. Eventually, however, the owner either has to reward the employee by permanently giving her more discretion than she would if she were not obligated by her past promises, or she has to punish him by giving him permanently less discretion. In either case, the owner is no longer able to make efficient use of the employee's information, the firm's decision making becomes inertial, and its performance declines. We then showed that our model speaks to the failure of established firms to adapt to changes in their environments, even when those changes are publicly observed. And we showed that it provides a rationale for persistent differences in the organization and performance of seemingly similar firms that various studies have documented over recent years.

The empirical literature on the allocation of control within firms is very small and, to our knowledge, there are no papers that explore how this allocation evolves over time. There is, however, some recent empirical evidence which does suggest that dynamic considerations may be playing an important role in the allocation of control. In particular, Bloom et al. (2012) provide evidence for the intuitive observation that trust facilitates delegation and, through this channel, can have a positive impact on firm performance. Even though their data is cross-sectional, they acknowledge at the end of their paper that trust is an inherently dynamic issue: "*we have considered trust as being exogenously endowed on firms and countries due to long-run effects of history and culture (such as religion). But corporate cultures do change over time, and modeling the endogenous evolution of trust and incentives to invest in it would be a fascinating avenue for future research.*"

Since we model delegation explicitly as a trust game, our model allows us to take a first step in this direction. Viewed through this lens, our model suggests that trust and delegation are interdependent and that they may evolve over time in ways that are perhaps counter-intuitive, at least at first. To see this, note that under the optimal relational contract, the owner starts out trusting her employee to act in her interest. Eventually, however, the owner has to either reward the employee by giving him more discretion or punish him by taking away his discretion entirely. In either case, the owner no longer trusts the employee to take her interests into account. Trust therefore unambiguously declines over time. Yet, the employee's discretion actually increases whenever the owner has to reward him for his performance early on. These dynamics contrast with the tempting view that since trust develops gradually and facilitates delegation, both should increase over time. And they suggest that the interdependence between trust and delegation is less obvious than it may at first appear and thus benefit from further, careful examination.

Appendix A: (Incomplete)

8.1 Additional Formal Results

LEMMA 6: *There exist four cutoffs* $\underline{u}_{CD}, \bar{u}_{CD}, \underline{u}_N, \bar{u}_N$ *such that the PPE payoff frontier* $\pi_2(u)$ *is divided into six regions:*

(i.) For $u \in [0,a]$, $\pi_2(u) = u$ *and* $(u, \pi_2(u))$ *is supported by randomization between exit and centralization.*

(ii.) For $u \in [a, \underline{u}_{CD}]$, $\pi_2(u) = ((\underline{u}_{CD} - u)\, a + (u-a)\, \pi_2(\underline{u}_{CD})) / (\underline{u}_{CD} - a)$ *and* $(u, \pi_2(u))$ *is supported by randomization between centralization and cooperative delegation.*

(iii.) For $u \in [\underline{u}_{CD}, \underline{u}_N]$, $\pi_2(u) = \pi_{2,D_C}(u)$ *and* $(u, \pi_2(u))$ *is supported by cooperative delegation.*

(iv − a.) For $u \in [\underline{u}_N, \tilde{u}_N]$, $\pi_2(u) = ((\tilde{u}_N - u)\, \pi_2(\underline{u}_N) + (u - \underline{u}_N)\, \tilde{\pi}_N) / (\tilde{u}_N - \underline{u}_N)$ *and* $(u, \pi_2(u))$ *is supported by a randomization between cooperative delegation and the new project.*

(iv − b.) For $u \in [\tilde{u}_N, \bar{u}_N]$, $\pi_2(u) = ((\bar{u}_N - u)\, \tilde{\pi}_N + (u - \tilde{u}_N)\, \pi_2(\bar{u}_N)) / (\bar{u}_N - \tilde{u}_N)$ *and* $(u, \pi_2(u))$ *is supported by a randomization between cooperative delegation and the new project.*

(v.) For $u \in [\bar{u}_N, \bar{u}_{CD}]$, $\pi_2(u) = \pi_{2,D_C}(u)$ *and* $(u, \pi_2(u))$ *is supported by cooperative delegation.*

(vi.) For $u \in [\bar{u}_{CD}, B]$, $\pi_2(u) = ((B-u)\, \pi_2(\bar{u}_{CD}) + (u - \underline{u}_{CD})\, b) / (B - \underline{u}_{CD})$ *and* $(u, \pi_2(u))$ *is supported by a randomization between cooperative and uncooperative delegation.*

LEMMA 7: *The PPE frontier* $\pi(u)$ *can be divided into five regions.*

$$\pi(u) = \begin{cases} u\pi(\underline{u}_{C_D})/\underline{u}_{C_D} & u \in [0, \underline{u}_{C_D}); \\ \pi_{C_D}(u) & u \in \left[\underline{u}_{C_D}, \bar{u}_{C_D}\right]; \\ \left(\left(\underline{u}_{D_C} - u\right)\pi(\bar{u}_{C_D}) + (u - \bar{u}_{C_D})\,\pi(\underline{u}_{D_C})\right) / \left(\underline{u}_{D_C} - \bar{u}_{C_D}\right) & u \in (\bar{u}_{C_D}, \underline{u}_{D_C}); \\ \pi_{D_C}(u) & u \in \left[\underline{u}_{D_C}, \bar{u}_{D_C}\right]; \\ ((B-u)\,\pi(\bar{u}_{D_C}) + (u - \bar{u}_{D_C})\, b) / (B - \bar{u}_{D_C}) & u \in (\bar{u}_{D_C}, B], \end{cases}$$

where $\underline{u}_{C_D} \geq (1-\delta)\, a$ *and* $\bar{u}_{D_C} = (1-\delta)\, b + \delta B$.

8.2 Proofs

LEMMA 1. *The PPE payoff set* $\mathcal{E}$ *has the following properties: (i.) it is compact, (ii.)* $\pi(u)$ *is concave, (iii.) the payoff pair* (u, π) *belongs to* $\mathcal{E}$ *if and only if* $u \in [0, B]$ *and* $\pi \in [bu/B, \pi(u)]$.

Proof of Lemma 1: 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 both $(0,0)$ and (B,b) are PPE payoffs sustained by termination and entrenchment respectively. In addition, there exists no actions that give the the agent payoffs below 0 or above B. This implies that if $(u,\pi) \in \mathcal{E}$, then $u \in [0,B]$. Moreover, one must have $\pi \leq bu/B$ because no payoff below the line segment joining $(0,0)$ and (B,b) is feasible. Next, given the public randomization device, any payoff on the line segment between $(0,0)$ and (B,b) can be supported as a PPE payoff. In other words, $(u, bu/B)$ is a PPE payoff for any $u \in [0,B]$. Finally, the randomization between $(u, bu/B)$ and $(u, \pi(u))$ allows us to obtain any payoff (u,π) for all $\pi \in [bu/B, \pi(u)]$. ■

LEMMA 2: *For any payoff $(u,\pi(u))$ on the frontier, the equilibrium continuation payoffs remain on the frontier.*

Proof of Lemma 2 : 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 control. 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 control but in which the continuation payoff is given by $(u_C, \widehat{\pi}_C)$, where $\widehat{\pi}_C = \pi_C + \varepsilon$ and where $\varepsilon > 0$ is small enough such that $\pi_C + \varepsilon \leq \pi(u_C)$. It follows from the promise-keeping constraints $\mathrm{PK}_{\mathrm{C}}^{\mathrm{P}}$ and $\mathrm{PK}_{\mathrm{C}}^{\mathrm{A}}$ that under this alternative strategy profile the payoffs are given by $\widehat{u} = u$ and $\widehat{\pi}_C = \pi(u) + \delta\varepsilon > \pi(u)$. It can be checked that this alternative strategy profile satisfies all the constraints and therefore constitutes a PPE. Since $\widehat{\pi}_C > \pi(u)$, 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 cooperative or uncooperative delegation.

■

LEMMA 3: In addition, the following hold.

(i.) If $(u,\pi(u))$ is supported with cooperative delegation, the Subordinate's continuation payoff can be chosen by

$$u_h(u) = \frac{u - (1-\delta)b}{\delta};$$
$$u_\ell(u) = \frac{u - (1-\delta)B}{\delta}.$$

(ii.) If $(u,\pi(u))$ is supported with uncooperative delegation, the Subordinate's continuation payoff is given by

$$u_{D_U}(u) = \frac{u - (1-\delta)B}{\delta}.$$

(iii.) If $(u,\pi(u))$ *is supported with control, the Subordinate's continuation payoff is given by*

$$u_C(u) = \frac{u-(1-\delta)a}{\delta}.$$

Proof of Lemma 3: For part $(i.)$, let $(u,\pi(u))$ be associated with the continuation payoffs $(u_\ell, u_h, \pi(u_\ell), \pi(u_h))$. Suppose that for this PPE the IC_{D_C} is slack, that is, $(1-\delta)b+\delta u_h > (1-\delta)B+\delta u_\ell$. Now consider an alternative strategy profile with continuation payoffs given by $(\widehat{u}_\ell, \widehat{u}_h, \pi(\widehat{u}_\ell), \pi(\widehat{u}_h))$, where $\widehat{u}_\ell = u_\ell + p\varepsilon$ and $\widehat{u}_h = u_h - (1-p)\varepsilon$ for $\varepsilon > 0$. It follows from the promise-keeping constraints $\mathrm{PK}^{\mathrm{P}}_{D_C}$ and $\mathrm{PK}^{\mathrm{A}}_{D_C}$ that, under this strategy profile, the payoffs are given by $\widehat{u} = u$ and

$$\widehat{\pi} = p\left[(1-\delta)B + \delta\pi(\widehat{u}_h)\right] + (1-p)\left[(1-\delta)b + \delta\pi(\widehat{u}_l)\right].$$

From the concavity of π it then follows that

$$\widehat{\pi} \geq (1-\delta)b + \delta((1-p)\pi(u_\ell) + p\pi(u_h)) = \pi(u).$$

It can be checked that for sufficiently small ε this alternative strategy profile satisfies all the constraints and therefore constitutes a PPE. Since $\widehat{\pi} \geq \pi(u)$ this implies that for any PPE with payoffs $(\pi, u(\pi))$ for which IC is not binding there exists another PPE for which $\mathrm{IC}_{\mathrm{D_C}}$ is binding and which gives the parties weakly larger payoffs. Notice that when $\mathrm{IC}_{\mathrm{D_C}}$ is binding, we have $u_h(u) = (u-(1-\delta)b)/\delta$ and $u_\ell(u) = (u-(1-\delta)B)/\delta$. This proves part $(i.)$. Parts $(ii.)$ and $(iii.)$ follow directly from the promise-keeping constraints $\mathrm{PK}^{\mathrm{A}}_{\mathrm{D_U}}$ and $\mathrm{PK}^{\mathrm{A}}_{\mathrm{C}}$.■

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

$$\pi(u) = \max_{q_j \geq 0, u_j \in [0,B]} \sum_{j \in \{C, D_C, D_U, E\}} q_j \pi_j(u_j)$$

such that

$$\sum_{j \in \{C, D_C, D_U, E\}} q_j = 1$$

and

$$\sum_{j \in \{C, D_C, D_U, E\}} q_j u_j = u.$$

Proof of Lemma 4: Since the frontier is Pareto efficient, by 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, even if the maximizers are not unique, the maximum is.

Next, instead of proving Lemma 5 directly, we establish alternative Lemma 5A-5C, which leads to Lemma 5$'$. Lemma 5$'$ provides a complete characterization of the PPE payoff frontier for all $p \in (0,1)$, which includes the case of $p \leq 1/2$ (Lemma 5).

LEMMA 5A: *There exists a cutoff value* $\bar{u}_{CD} = (1-\delta)\, b + \delta B$ *such that* $\pi_{D_U}(u) = \pi(u)$ *if and only if*

$$u \in [(1-\delta)\, B + \delta \bar{u}_{CD}, B].$$

Proof of Lemma 5A: First, notice that $\pi_{D_U}(B) = \pi(B)$. Next, recall that $\pi_{D_U}(u) = (1-\delta)\, b + \delta\pi(u_{D_U}(u))$. Taking the right derivative, we have

$$\pi'^{+}_{D_U}(u) = \delta \pi'^{+}(u_{D_U}(u))\, u'^{+}_{D_U}(u) = \pi'^{+}(u_{D_U}(u)) \geq \pi'^{+}(u),$$

where we used the fact that if $u < B$, $u_{UD}(u) < u$, and therefore $\pi'^{+}(u_{D_U}(u)) \geq \pi'^{+}(u)$ by concavity of the frontier. Since $\pi'^{+}_{D_U}(u) \geq \pi'^{+}(u)$ for all $u < B$, there exists u^* such that $\pi_{D_U}(u) = \pi(u)$ if and only if $u \in [u^*, B]$. Next, we show that $u^* = (1-\delta)\, B + \delta \bar{u}_{CD}$.

To do this, we first show that there exists some $u < B$ such that $\pi_{D_U}(u) = \pi(u)$, i.e., $u^* < B$. We prove this by contradiction. Suppose to the contrary that $\pi_{D_U}(u) < \pi(u)$ 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 D_U, and if it were supported by C or D_C, then the agent's continuation payoffs (u_C and u_h) 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. Let 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. $u_{UD}(u) \geq B - \varepsilon$), we have

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

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

The above shows that $\pi_{D_U}(u) = \pi(u)$ for $u \in [u^*, B]$, where $u^* < B$. It follows that for all $u \in [u^*, B]$, $\pi'^{+}_{D_U}(u) = \pi'^{+}(u_{D_U}(u)) = \pi'^{+}(u)$. Since π is concave, this implies that the slope of π is constant for all Subordinate's payoffs in $(u_{D_U}(u), u)$. It is then immediate that $\pi(u)$ is a straight line between $[u^{-1}_{D_U}(u^*), B]$. Let $(u', \pi(u'))$ be the left end point of the line segment. Notice that $(u', \pi(u'))$ is an extremal point of the payoff frontier, it must be supported by pure action. Moreover, it cannot be supported by uncooperative delegation, because $u' \leq u^{-1}_{D_U}(u^*) < u^*$. This implies that $(u', \pi(u'))$ is supported by either cooperative delegation or control. Next, we show that $(u', \pi(u'))$ cannot be supported by control, so we must have $\pi(u') = \pi_{D_C}(u')$.

Suppose to the contrary that $\pi(u') = \pi_C(u')$. Notice that

$$\begin{aligned} u' &= (1-\delta)a + \delta u_C(u'); \\ \pi(u') &= (1-\delta)a + \delta\pi(u_C(u')), \end{aligned}$$

which implies that

$$\begin{aligned} u' - a &= \delta(u_C(u') - a); \\ \pi(u') - a &= \delta(\pi(u_C(u')) - a). \end{aligned}$$

Now take any $u \in (\max\{a, u'\}, B)$. We have $u_C(u) > u$, and the above implies that

$$\frac{\pi(u) - a}{u - a} = \frac{\pi(u_C(u)) - a}{u_C(u) - a} = \frac{\pi(u_C(u)) - \pi(u')}{u_C(u) - u} = \frac{\pi(u) - b}{u - B},$$

where the last inequality holds, because π is a straight line to the right of u' and $u > u'$. The equalities then imply that $(u, \pi(u))$ lies on the line segment between (a, a) and (B, b) for all $u \in [a, B]$, which contradicts [NOTE: Assumption X] ($B - a \leq \frac{(1-\delta)}{\delta}(B - b)$). This implies that $\pi(u') > \pi_C(u')$, and we must have $\pi(u') = \pi_{D_C}(u')$.

Finally, we show that $u' = \bar{u}_{CD}$, i.e., $u_h(u') = B$. By $\mathrm{SE_{CD}}$, the continuation payoff $u_h(u')$ satisfies $u_h(u') \leq B$. Now suppose to the contrary that $u_h(u') < B$. Recall that s is the slope of the payoff frontier between $(u', \pi(u'))$ and (B, b). Now consider an alternative strategy profile that is supported by cooperative delegation and whose continuation payoffs are given by $(\widehat{u}_\ell, \widehat{u}_h, \pi(\widehat{u}_\ell), \pi(\widehat{u}_h))$, where $\widehat{u}_\ell = u_\ell(u') + \varepsilon$ and $\widehat{u}_h = u_h(u') + \varepsilon$ for small $\varepsilon > 0$. It follows from the promise-keeping constraints $\mathrm{PK_{CD}^P}$ and $\mathrm{PK_{CD}^A}$ that, under this strategy profile, the payoffs are given by $\widehat{u} = u' + \delta\varepsilon$ and

$$\begin{aligned} \widehat{\pi} &= p[(1-\delta)B + \delta\pi(\widehat{u}_h)] + (1-p)[(1-\delta)b + \delta\pi(\widehat{u}_\ell)] \\ &= \pi(u') + p\delta[\pi(u_h(u') + \varepsilon) - \pi(u_h(u'))] \\ &\quad + (1-p)\delta[\pi(u_\ell(u') + \varepsilon) - \pi(u_\ell(u'))] \\ &> \pi(u') + \delta s\varepsilon. \end{aligned}$$

Notice that the strict inequality follows, because s is the smallest slope of π, and since $u_\ell(u') < u'$, we have $\pi(u_\ell(u') + \varepsilon) - \pi(u_\ell(u')) > s\varepsilon$ by the definition of u'. It can be checked that for sufficiently small ε this alternative strategy profile satisfies all the constraints and therefore constitutes a PPE. This implies that

$$\pi\left(u'\right)+\delta s\varepsilon<\widehat{\pi}\leq\pi\left(u'+\delta\varepsilon\right)=\pi(u')+\delta s\varepsilon,$$

where the weak inequality follows from the definition of π and the equality follows from that π is a straight line with slope s to the right of u'. Since the above chain of inequalities lead to a contradiction, we must have $u_h(u')=B$, or equivalently, $u'=\bar{u}_{CD}$. ■

LEMMA 5B: *There exists a cutoff probability level $p^{**}>1/2$ such that the following hold.*

(i.) If $p\in(p^{**},1)$, $\pi_C(u)<\pi(u)$ for all $u\in[0,B]$.

(ii.) If $p\in(0,p^{**}]$, there exists $\underline{u}_C\leq a\leq\overline{u}_C$ such that $\pi_C(u)=\pi(u)$ if and only if $u\in[(1-\delta)\,a+\delta\underline{u}_C,(1-\delta)\,a+\delta\overline{u}_C]$. In addition, $\underline{u}_C=0$ if and only if $p\leq1/2$.

Proof of Lemma 5B: First, we show that if $\pi(a)<a$, then $\pi_C(u)<\pi(u)$ for all $u\in[0,B]$.

To prove this, we show that if that $\pi_C(u)=\pi(u)$ for some $u\neq a$, then $\pi(a)=a$. Now suppose there exists $u>a$ such that $\pi_C(u)=\pi(u)$. Since $u>a$, we have $u_C\left(u\right)=\left(u-\left(1-\delta\right)a\right)/\delta>u$. Recall that $\pi_C\left(u\right)=\left(1-\delta\right)a+\delta\pi\left(u_C\left(u\right)\right)$. Taking left derivatives, we have

$$\pi_C^{'-}\left(u\right)=\pi^{'-}\left(u_C\left(u\right)\right)\leq\pi^{'-}\left(u\right),$$

where the inequality follows because the payoff frontier is concave and $u_C\left(u\right)>u$. Since π is concave, the above then implies that π is a straight line in $\left[u_C\left(u\right),u\right]$, and $\pi_C\left(u'\right)=\pi\left(u'\right)$ for all $u'\in\left[u_C\left(u\right),u\right]$. Repeating the argument at $u_C\left(u\right)$, we get that $\pi_C\left(u''\right)=\pi\left(u''\right)$ for all $u''\in[u_C(u_C\left(u\right)),u_C\left(u\right)]$. Using the argument repeatedly and using the fact that π is continuous at a, we then have that $\pi_C\left(u'\right)=\pi\left(u'\right)$ for all $u'\in[a,u]$. It follows that

$$\pi\left(a\right)=\pi_C\left(a\right)=\left(1-\delta\right)a+\delta\pi\left(u_C\left(a\right)\right)=\left(1-\delta\right)a+\delta\pi\left(a\right),$$

which implies that $\pi(a)=a$, contradicting the assumption that $\pi(a)<a$. This proves that if $\pi(a)<a$, then $\pi_C(u)<\pi(u)$ for all $u>a$. The identical argument as above can be used to show that if $\pi(a)<a$, then $\pi_C(u)<\pi(u)$ for all $u<a$. This finishes showing that if $\pi(a)<a$, then $\pi_C(u)<\pi(u)$ for all $u\in[0,B]$.

Next, when $\pi(a)=a$, define $\overline{u}_C=\max\{u:\pi_C(u_C(u))=\pi(u_C(u))\}$. Notice that the argument above implies that $\pi'(u)$ is the same for all u between a and $\overline{u}_C$, and this shows that π is a line segment in $[a,\overline{u}_C]$. Similarly, define $\underline{u}_C=\min\{u:\pi_C(u_C(u))=\pi(u_C(u))\}$, and we then have that π is a line segment in $[\underline{u}_C,a]$.

Notice that when $p \leq 1/2$, the set of feasible payoffs lie below the 45 degree line for $u \in [0, a]$. In addition, payoffs on the 45 degree line with end points in $(0, 0)$ and (a, a) can be implemented by the randomization between the two end points (that are both PPE payoffs). This implies that $\pi(u) = u$ for all $u \in [0, a]$ and immediately implies that $\underline{u}_C = 0$.

Finally, we will show in the proof of next lemma that there exists $p^{**} > 1/2$ such that $\pi(a) > a$ if and only $p > p^{**}$. Given this result, Lemma 5B follows immediately. We prove this result in Lemma 5C for convenience, and its proof does not depend on results established in this lemma. ■

LEMMA 5C: *There exists a cutoff probability level $p^{**} > 1/2$ such that the following holds.*

(i.) If $p \in [p^{**}, 1)$, $\pi_{D_C}(u) = \pi(u)$ if and only if $u \in [\underline{u}_{CD}, \bar{u}_{CD}]$, where $u_1 = (1-\delta) B$.

(ii.) If $p \in [1/2, p^{**}]$, $\pi_{D_C}(u) = \pi(u)$ if and only if $u \in [\underline{u}^A_{CD}, \bar{u}^A_{CD}] \cup [\underline{u}_{CD}, \bar{u}_{CD}]$ for some $\bar{u}^A_{CD} < a < \underline{u}_{CD} \leq \delta a + (1-\delta) B$.

(iii.) If $p \in (0, 1/2)$, $\pi_{D_C}(u) = \pi(u)$ if and only if $u \in [\underline{u}_{CD}, \bar{u}_{CD}]$ for some $\underline{u}_{CD} \in (a, \delta a + (1-\delta) B]$.

Proof of Lemma 5C: To prove the lemma, we take the following steps.

Step 1: We establish properties of PPE payoff of a modified game.

Consider a modified game in which C is not feasible. Let π_T be the associated PPE payoff frontier of the modified game. We estabilish the following properties of π_T.

A: For $u \in [\bar{u}_{CD}, B]$, where $\bar{u}_{CD} = (1-\delta) b + \delta B$,

$$\pi_T(u) = \frac{u - \bar{u}_{CD}}{B - \bar{u}_{CD}} b + \frac{B - u}{B - \bar{u}_{CD}} \pi_T(\bar{u}_{CD}).$$

B: For $u \in [0, \underline{u}_{CD}]$, where $\underline{u}_{CD} = (1-\delta) B$,

$$\pi_T(u) = \frac{u}{\underline{u}_{CD}} \pi_T(\underline{u}_{CD}).$$

C: For $u \in [\underline{u}_{CD}, \bar{u}_{CD}]$,

$$\pi_T(u) = p\left[(1-\delta) B + \delta \pi_T(u_h(u))\right] + (1-p)\left[(1-\delta) b + \delta \pi_T(u_\ell(u))\right].$$

The properties of π_T are established in the same way as the method used here except it is simpler. We therefore omit the proof of Step 1.

Step 2: There exists $p^{**} > 1/2$ such that $\pi_T(u, p) = \pi(u, p)$ if and only if $p \geq p^{**}$.

To see this, we first show that for all $u \in (0, B)$, $\pi_T(u, p)$ is strictly increasing in p. Consider $p_1 > p_2$. Define the operator $T_i f$, which maps bounded nonnegative functions on $[0, B]$ to bounded nonnegative functions on $[0, B]$, as

$$
\begin{aligned}
T_i f(u) &= \frac{u}{\underline{u}_{CD}} f(\underline{u}_{CD}) \text{ for } u \leq \underline{u}_{CD} \\
T_i f(u) &= \frac{u - \bar{u}_{CD}}{B - \bar{u}_{CD}} b + \frac{B - u}{B - \bar{u}_{CD}} f(\bar{u}_{CD}) \text{ for } \bar{u}_{CD} \leq u \leq B \\
T_i f(u) &= p_i \left[(1 - \delta) B + \delta f(u_h(u))\right] + (1 - p_i) \left[(1 - \delta) b + \delta f(u_\ell(u))\right] \text{ for } \underline{u}_{CD} \leq u \leq \bar{u}_{CD}
\end{aligned}
$$

It is clear that T_i is bounded monotone in the sense that $T_i(f_1) \geq T_i(f_2)$ whenever $f_1 \geq f_2$ (in the sense that $f_1(u) \geq f_2(u)$ for all u). Let $Z(u) \equiv 0$ on $[0, B]$, it follows that $Z^* \equiv \lim_{n\to\infty} T_i(z)$ is a fixed point of T_i. Moreover, the fixed point is unique. Suppose to the contrary that f_1 and f_2 are two fixed point of T_i. Let $M = \sup_{u\in[0,B]}\{|f_1(u) - f_2(u)|\}$. Now notice that

$$
\begin{aligned}
M &= \sup_{u\in[\underline{u}_{CD}, \bar{u}_{CD}]} \{|f_1(u) - f_2(u)|\} \\
&= \sup_{u\in[\underline{u}_{CD}, \bar{u}_{CD}]} \{|Tf_1(u) - Tf_2(u)|\} \\
&= \sup_{u\in[\underline{u}_{CD}, \bar{u}_{CD}]} \{\delta |p_i (f_1(u_h(u)) - f_2(u_h(u))) + (1 - p_i)(f_1(u_l(u)) - f_2(u_l(u)))|\} \\
&\leq \delta p_i \sup_{u\in[0,B]} \{|f_1(u) - f_2(u)|\} + \delta(1 - p_i) \sup_{u\in[0,B]} \{|f_1(u) - f_2(u)|\}. \\
&\leq \delta M,
\end{aligned}
$$

and this implies that $M = 0$. This shows that T_i has a unique fixed point.

Now let $\pi_T(u, p_1)$ be the unique fixed point of T_1 and $\pi_T(u, p_2)$ be the unique fixed point of T_2. Notice that for each $u \in [\underline{u}_{CD}, \bar{u}_{CD}]$,

$$
\begin{aligned}
T_1 \pi_T(u, p_2) &= p_1 \left[(1 - \delta) B + \delta \pi_T(u_h(u), p_2)\right] + (1 - p_1) \left[(1 - \delta) b + \delta \pi_T(u_h(u), p_2)\right] \\
&= (p_1 - p_2)(1 - \delta)(B - b) + T_2 \pi_T(u, p_2) \\
&> \pi_T(u, p_2).
\end{aligned}
$$

The monotonicity of T_1 then implies immediately that $T_1 \pi_T(u, p_2) > \pi_T(u, p_2)$ for all $u \in (0, B)$. This finishes showing that for all $u \in (0, B)$, $\pi_T(u, p)$ is strictly increasing in p.

Next, we show that there exists p^{**} such that $\pi_T(u, p) = \pi(u, p)$ if and only if $p \geq p^*$. Choose p^{**} as the cutoff value such that $\pi_T(a, p^{**}) = a$. This implies that when $p > p^{**}$, $a < \pi_T(a, p) \leq \pi(a, p)$. By the argument in Lemma 5B, we then have that the frontier $\pi(u, p)$ is not supported by centralization for all $u \in [0, B]$, and this implies that $\pi_T(u, p) = \pi(u, p)$ for all $u \in [0, B]$. This implies that $\pi_T(u,p) = \pi(u,p)$ for all $p > p^{**}$, and since both π_T and π are continuous in p (by the maximum theorem), we then also have $\pi_T(u,p) = \pi(u,p)$.

Notice that we must have $p^{**} \geq 1/2$ because otherwise (a,a) lies strictly above the set of feasible payoffs and thus cannot be a PPE payoff. When $p = 1/2$, let u^* be the maximal equilibrium payoff of the agent such that $\pi_T(u)$ is on the 45 degree line. If $u^* > 0$, then it satisfies $u_h(u^*) = u^*$, which implies that $u^* = b < a$. Consequently, (a,a) again lies above $(a, \pi_T(a))$ when $p = 1/2$. As a result, we have $p^{**} > 1/2$. By combining Step 1 and 2, this proves part (i) of Lemma 5C.

Step 3: When $p < 1/2$, then $\pi_{D_C}(u) = \pi(u)$ if and only if $u \in [\underline{u}_{CD}, \bar{u}_{CD}]$ for some $\underline{u}_{CD} \in (a, \delta a + (1-\delta)B]$.

To see this, notice that for $u \in [\overline{u}_C, \bar{u}_{CD}]$, where recall that $\overline{u}_C = \max\{u : \pi_C(u_C(u)) = \pi(u_C(u))\}$, $\pi(u)$ cannot be supported by centralization or uncooperative delegation by Step 2 and Lemma 5A. Therefore, for any $u \in [\overline{u}_C, \bar{u}_{CD}]$, either $\pi(u) = \pi_D(u)$ (in which case we are done) or there exists a $\rho \in (0,1)$, a $\widehat{u}_1 \in [\overline{u}_C, u)$, and a $\widehat{u}_2 \in (u, \bar{u}_{CD}]$ such that (i.) both $(\widehat{u}_1, \pi(\widehat{u}_1))$ and $(\widehat{u}_2, \pi(\widehat{u}_2))$ satisfy $\pi_{D_C}(\widehat{u}_i) = \pi(\widehat{u}_i)$ for $i = 1, 2$, (ii.) $(u, \pi(u)) = \rho(\widehat{u}_1, \pi(\widehat{u}_1)) + (1-\rho)(\widehat{u}_2, \pi(\widehat{u}_2))$. Let $\widehat{u}_{ih}$ and $\widehat{u}_{il}$, $i = 1, 2$ be the agent's associated continuation payoffs for the two PPE.

Now consider an alternative strategy profile in which cooperative delegation is chosen and the continuation payoff given by $(\widehat{u}_h, \widehat{u}_l, \pi(\widehat{u}_h), \pi(\widehat{u}_l))$, where $\widehat{u}_h = \rho\widehat{u}_{1h} + (1-\rho)\widehat{u}_{2h}$ and $\widehat{u}_l = \rho\widehat{u}_{1l} + (1-\rho)\widehat{u}_{2l}$. It follows from the promise keeping constraints PK_{PD} and PK_{AD} that under this strategy profile the payoffs are given by $\widehat{u} = u$ and

$$\begin{aligned}
\widehat{\pi} &= p\left[(1-\delta)B + \delta\pi\left(\rho\widehat{u}_{1h} + (1-\rho)\widehat{u}_{2h}\right)\right] \\
&\quad + (1-p)\left[(1-\delta)b + \delta\pi\left(\rho\widehat{u}_{1h} + (1-\rho)\widehat{u}_{2h}\right)\right] \\
&\geq \rho\pi(\widehat{u}_1) + (1-\rho)\pi(\widehat{u}_2) \\
&= \pi(u).
\end{aligned}$$

It can be checked that this alternative strategy profile satisfies all the constraints and therefore constitutes a PPE. Therefore, $\pi_{D_C}(u) \geq \widehat{\pi} \geq \pi(u)$. This proves that $\pi_{D_C}(u) = \pi(u)$ if $u \in [\overline{u}_C, \bar{u}_{CD}]$.

Next, define $\underline{u}_{CD} = \min\{u : \pi_{D_C}(u) = \pi(u)\}$. By the above, we see that $\underline{u}_{CD} \leq \overline{u}_C$. In addition, since $p < 1/2$, it is clear that $\pi_{D_C}(u) < u = \pi(u)$ for all $u \in [0, a]$, where the equality follows from the proof of Lemma 5B. Therefore, $\underline{u}_{CD} > a$. It remains to show that $\pi_{D_C}(u) = \pi(u)$ if $u \in [\underline{u}_{CD}, \overline{u}_C]$. Suppose $u = \rho\underline{u}_{CD} + (1-\rho)\overline{u}_C$. Consider a strategy profile in which cooperative delegation is chosen and the continuation payoff given by $(\widehat{u}_h, \widehat{u}_l, \pi(\widehat{u}_h), \pi(\widehat{u}_l))$, where $\widehat{u}_h = \rho u_h(\underline{u}_{CD}) + (1-\rho)u_h(\overline{u}_C)$

and $\widehat{u}_l = \rho u_l(\underline{u}_{CD}) + (1-\rho)u_l(\overline{u}_C)$. It follows from the promise keeping constraints $\mathrm{PK_{PD}}$ and $\mathrm{PK_{AD}}$ that under this strategy profile the payoffs are given by $\widehat{u} = u$ and

$$
\begin{aligned}
\widehat{\pi} &= p\left[(1-\delta)B + \delta\pi\left(\rho u_h(\underline{u}_{CD}) + (1-\rho)u_h(\overline{u}_C)\right)\right] \\
&\quad + (1-p)\left[(1-\delta)b + \delta\pi\left(\rho u_l(\underline{u}_{CD}) + (1-\rho)u_l(\overline{u}_C)\right)\right] \\
&\geq \rho\pi\left(\underline{u}_{CD}\right) + (1-\rho)\pi\left(\overline{u}_C\right) \\
&= \pi(u),
\end{aligned}
$$

where the last equality follows from the linearity of π in $[\underline{u}_{CD}, \overline{u}_C]$. Therefore, $\pi_{D_C}(u) \geq \widehat{\pi} \geq \pi(u)$, and, thus, $\pi_{D_C}(u) = \pi(u)$. This proves that $\pi_{D_C}(u) = \pi(u)$ for all $u \in [\underline{u}_{CD}, \bar{u}_{CD}]$.

Next, we prove $\underline{u}_{CD} \leq \delta a + (1-\delta)B$, or equivalently, $u_l(\underline{u}_{CD}) \leq a$, and this finishes the proof of Step 3. Now suppose to the contrary that $u_l(\underline{u}_{CD}) > a$. This implies that $u_l(\overline{u}_C) > a$ since $\overline{u}_C \geq \underline{u}_{CD}$. Recall (in Lemma 5B) that π is a straight line in $[a, \overline{u}_C]$. Denote its slope as s. Now consider a strategy profile that is supported by delegation and whose continuation payoffs are given by $(\widehat{u}_l, \widehat{u}_h, \pi(\widehat{u}_l), \pi(\widehat{u}_h))$, where $\widehat{u}_l = u_l(\overline{u}_C) - \varepsilon$ and $\widehat{u}_h = u_h(\overline{u}_C) - \varepsilon$ for $\varepsilon > 0$. We choose ε small enough so that $u_l(\overline{u}_C) - \varepsilon > a$ and $u_h(\overline{u}_C) - \varepsilon > \overline{u}_C$. It follows from the promise keeping constraints $\mathrm{PK_{PD}}$ and $\mathrm{PK_{AD}}$ that under this strategy profile the payoffs are given by $\widehat{u} = \overline{u}_C - \delta\varepsilon$ and

$$
\begin{aligned}
\widehat{\pi} &= p\left[(1-\delta)B + \delta\pi\left(\widehat{u}_h\right)\right] + (1-p)\left[(1-\delta)b + \delta\pi\left(\widehat{u}_l\right)\right] \\
&= \pi\left(\overline{u}_C\right) + p\delta\left[\pi\left(u_h\left(\overline{u}_C\right) - \varepsilon\right) - \pi\left(u_h\left(\overline{u}_C\right)\right)\right] \\
&\quad + (1-p)\,\delta\left[\pi\left(u_\ell\left(\overline{u}_C\right) - \varepsilon\right) - \pi\left(u_\ell\left(\overline{u}_C\right)\right)\right] \\
&> \pi\left(\overline{u}_C\right) - \delta s\varepsilon.
\end{aligned}
$$

Notice that the strict inequality follows because $\pi\left(u_\ell\left(\overline{u}_C\right) - \varepsilon\right) - \pi\left(u_\ell\left(\overline{u}_C\right)\right) = -s\varepsilon$ (since $u_\ell\left(\overline{u}_C\right) - \varepsilon > a$) and $\pi\left(u_h\left(\overline{u}_C\right) - \varepsilon\right) - \pi\left(u_h\left(\overline{u}_C\right)\right) > -s\varepsilon$ since $u_h(\overline{u}_C) - \varepsilon > \overline{u}_C$ and $\pi'_-(u) < s$ for all $u > \overline{u}_C$. It can be checked that this alternative strategy profile satisfies all the constraints and therefore constitutes a PPE. This implies that

$$
\pi\left(\overline{u}_C\right) - \delta s\varepsilon < \widehat{\pi} \leq \pi\left(\overline{u}_C - \delta\varepsilon\right) = \pi(\overline{u}_C) - \delta s\varepsilon,
$$

where the weak inequality follows from the definition of π and the equality follows from that π is a straight line with slope s between a and $\overline{u}_C$. Since the above chain of inequalities lead to a contradiction, we must have $u_l(\underline{u}_{CD}) \leq u_l(\overline{u}_C) \leq a$, or equivalently $\underline{u}_{CD} \leq \delta a + (1-\delta)B$. This proves Step 3, and, thus, part (iii) of Lemma 5C.

Step 4: If $p \in [1/2, p^{*})$, $\pi_{D_C}(u) = \pi(u)$ if and only if $u \in [\underline{u}_{CD}^{A}, \bar{u}_{CD}^{A}] \cup [\underline{u}_{CD}, \bar{u}_{CD}]$ for some $\bar{u}_{CD}^{A} < a < \underline{u}_{CD} \leq \delta a + (1-\delta)B$, where $u_1 = (1-\delta)B$.

To prove this, notice that for $u \geq a$, we can again define $\underline{u}_{CD} = \min\{u : \pi_D(u) = \pi(u), u \geq a\}$. Using the same argument as in Step 3, we can show that when $u > a$, $\pi_{D_C}(u) = \pi(u)$ if and only if $u \in [\underline{u}_{CD}, \bar{u}_{CD}]$. Next, notice that when $p \geq 1/2$, $\pi_{D_C}(u) \geq u$ for all u such that $u_l(u) \geq 0$ and $u_h(u) \leq a$, i.e., $u \in [(1-\delta)B, (1-\delta)b + \delta a]$. This implies that $\pi_{D_C}(u) = \pi(u)$ for some $u \in (0, a)$ (because otherwise $\pi(u) = u$ for all $u \in (0, a)$). Now define $\bar{u}_{CD}^{A} = \max\{u : \pi_{D_C}(u) = \pi(u), u \leq a\}$ and $\underline{u}_{CD}^{A} = \min\{u : \pi_D(u) = \pi(u)\}$. By the same argument as in Step 3, we can show that $\pi_{D_C}(u) = \pi(u)$ for all $u \in [\underline{u}_{CD}^{A}, \bar{u}_{CD}^{A}]$.

By the definition of $\underline{u}_{CD}^{A}$, it is clear that $\pi(u)$ is a straight line between 0 and $\underline{u}_{CD}^{A}$. Let the slope of this line segment be s. To see $\underline{u}_{CD}^{A} = (1-\delta)B$, notice that $\underline{u}_{CD}^{A} \geq (1-\delta)B$ because otherwise cooperative delegation is not feasible. Suppose to the contrary that $\underline{u}_{CD}^{A} > (1-\delta)B$. Consider a strategy profile that is supported by delegation and whose continuation payoffs are given by $(\widehat{u}_l, \widehat{u}_h, \pi(\widehat{u}_l), \pi(\widehat{u}_h))$, where $\widehat{u}_l = u_l(\underline{u}_{CD}^{A}) - \varepsilon$ and $\widehat{u}_h = u_h(\underline{u}_{CD}^{A}) - \varepsilon$ for $\varepsilon > 0$. We choose ε small enough so that $u_l(\underline{u}_{CD}^{A}) - \varepsilon > 0$. It follows from the promise keeping constraints $\mathrm{PK}_{\mathrm{PD}}$ and $\mathrm{PK}_{\mathrm{AD}}$ that under this strategy profile the payoffs are given by $\widehat{u} = \underline{u}_{CD}^{A} - \delta\varepsilon$ and

$$\begin{aligned}
\widehat{\pi} &= p\left[(1-\delta)B + \delta\pi(\widehat{u}_h)\right] + (1-p)\left[(1-\delta)b + \delta\pi(\widehat{u}_l)\right] \\
&= \pi\left(\underline{u}_{CD}^{A}\right) + p\delta\left[\pi\left(u_h\left(\underline{u}_{CD}^{A}\right) - \varepsilon\right) - \pi\left(u_h\left(\underline{u}_{CD}^{A}\right)\right)\right] \\
&\quad + (1-p)\delta\left[\pi\left(u_\ell\left(\underline{u}_{CD}^{A}\right) - \varepsilon\right) - \pi\left(u_\ell\left(\underline{u}_{CD}^{A}\right)\right)\right] \\
&\geq \pi\left(\underline{u}_{CD}^{A}\right) - \delta s\varepsilon.
\end{aligned}$$

Notice inequality follows because π is concave so s is its maximal derivative. It can be checked that this alternative strategy profile satisfies all the constraints and therefore constitutes a PPE. This implies that

$$\pi_{D_C}\left(\underline{u}_{CD}^{A} - \delta\varepsilon\right) \geq \widehat{\pi} \geq \pi\left(\underline{u}_{CD}^{A}\right) - \delta s\varepsilon = \pi\left(\underline{u}_{CD}^{A} - \delta\varepsilon\right),$$

where the equality follows from that π is a straight line with slope s between 0 and $\underline{u}_{CD}^{A}$. By the above chain of inequalities, $\pi_{D_C}\left(\underline{u}_{CD}^{A} - \delta\varepsilon\right) = \pi\left(\underline{u}_{CD}^{A} - \delta\varepsilon\right)$, which lead to a contradiction because $\underline{u}_{CD}^{A}$ is defined as the smallest agent's payoff such that $\pi_{D_C}(u) = \pi(u)$.

Next, notice that we must have $\bar{u}_{CD}^{A} < a$. Suppose to the contrary that $\bar{u}_{CD}^{A} = a$, then by definition we also have $\underline{u}_{CD} = a$, and this implies that $\pi(u) = \pi_{D_C}(u)$ for all $u \in [\underline{u}_{CD}^{A}, \bar{u}_{CD}]$. This implies that $\pi(u) = \pi_T(u)$, where recall that $\pi_T(u)$ is defined in Step 1 (where centralization is not used). This contradicts that assumption that $p < p^{**}$. Similarly, we must also have $\underline{u}_{CD} > a$. This proves Step 4, and, thus, part (ii) of Lemma 5C. ■

Combining Lemma 5A-5C, we obtain the following lemma.

Lemma 5': *The PPE frontier $\pi(u)$ satisfies the following.*

(i.) If $p \in [p^{**}, 1)$, the frontier can be divided into three regions.

$$\pi(u) = \begin{cases} u\pi(\underline{u}_{CD})/\underline{u}_{CD} & u \in [0, \underline{u}_{CD}); \\ \pi_{D_C}(u) & u \in [\underline{u}_{CD}, \bar{u}_{CD}]; \\ ((B-u)\,\pi(\bar{u}_{CD}) + (u-\bar{u}_{CD})\,b)\,/\,(B-\bar{u}_{CD}) & u \in (\bar{u}_{CD}, B], \end{cases}$$

where $\underline{u}_{CD} = (1-\delta)\,B$ and $\bar{u}_{CD} = (1-\delta)\,b + \delta B$.

(ii.) If $p \in (1/2, p^*)$, the frontier can be divided into six regions.

$$\pi(u) = \begin{cases} u\pi(\underline{u}^A_{CD})/\underline{u}^A_{CD} & u \in [0, \underline{u}^A_{CD}); \\ \pi_{D_C}(u) & u \in [\underline{u}^A_{CD}, \bar{u}^A_{CD}]; \\ \left((a-u)\,\pi(\bar{u}^A_{CD}) + \left(u-\bar{u}^A_{CD}\right)a\right)\,/\,\left(a-\bar{u}^A_{CD}\right) & u \in (\bar{u}^A_{CD}, a); \\ ((\underline{u}_{CD}-u)\,a + (u-a)\,\pi(\underline{u}_{CD}))\,/\,(\underline{u}_{CD}-a) & u \in [a, \underline{u}_{CD}); \\ \pi_{D_C}(u) & u \in [\underline{u}_{CD}, \bar{u}_{CD}]; \\ ((B-u)\,\pi(\bar{u}_{CD}) + (u-\bar{u}_{CD})\,b)\,/\,(B-\bar{u}_{CD}) & u \in (\bar{u}_{CD}, B], \end{cases}$$

where $\underline{u}^A_{CD} = (1-\delta)\,B$, $\underline{u}_{CD} \in (a, \delta a + (1-\delta)\,B)$ and $\bar{u}_{CD} = (1-\delta)\,b + \delta B$.

(iii.) If $p \in (0, 1/2]$, the frontier can be divided into four regions.

$$\pi(u) = \begin{cases} u & u \in [0, a); \\ ((\underline{u}_{CD}-u)\,a + (u-a)\,\pi(\underline{u}_{CD}))\,/\,(\underline{u}_{CD}-a) & u \in [a, \underline{u}_{CD}); \\ \pi_{D_C}(u) & u \in [\underline{u}_{CD}, \bar{u}_{CD}]; \\ ((B-u)\,\pi(\bar{u}_{CD}) + (u-\bar{u}_{CD})\,b)\,/\,(B-\bar{u}_{CD}) & u \in (\bar{u}_{CD}, B], \end{cases}$$

where $\underline{u}_{CD} \in (a, \delta a + (1-\delta)\,B)$ and $\bar{u}_{CD} = (1-\delta)\,b + \delta B$.

PROPOSITION 2: *In the optimal equilibrium, the principal chooses cooperative delegation if and only if $t \leq \tau$ for some random time period τ.* $\Pr(\tau < \infty) = 1$ *so that delegation occurs with probability 0 in the long run. In addition, there exists $p^* < p^{**}$ that the following holds.*

(i.) When $p \geq p^{**}$, either the relationship terminates or the agent entrenches in the long run:

$$\begin{aligned} \lim_{t\to\infty} \Pr(u_t &= 0) > 0,\ \lim_{t\to\infty} \Pr(u_t = B) > 0, \\ \lim_{t\to\infty} \Pr(u_t &= 0) + \Pr(u_t = B) = 1. \end{aligned}$$

(ii.) When $p \in (p^*, p^{**})$, the relationship terminates with positive probability less than 1. When the relationship does not terminate in the long run, principal chooses either centralization permanently or uncooperative delegation permanently.

$$\begin{aligned}\lim_{t\to\infty} \Pr(u_t &= 0) > 0,\ \lim_{t\to\infty} \Pr(u_t = a) > 0,\ \lim_{t\to\infty} \Pr(u_t = B) > 0,\\ \lim_{t\to\infty} \Pr(u_t &= 0) + \Pr(u_t = a) + \Pr(u_t = B) = 1.\end{aligned}$$

(iii.) When $p \leq p^*$, the principal chooses either centralization permanently or uncooperative delegation permanently in the long run:

$$\begin{aligned}\lim_{t\to\infty} \Pr(u_t &= a) > 0,\ \lim_{t\to\infty} \Pr(u_t = B) > 0,\\ \lim_{t\to\infty} \Pr(u_t &= a) + \Pr(u_t = B) = 1.\end{aligned}$$

Proof of Proposition 2: By Assumption Y, the relationship starts with cooperative delegation. In addition, if ever control, entrenchment, or termination is used, Lemma 5' immediately implies that the relationship stops there forever. This establishes the existence of the random time. Moreover, since these are the only absorbing states of the relationship, and it can be easily show that there exists an $\varepsilon > 0$ and a large enough N such that the probability the relationship ends in one of the absorbing state exceeds ε every N periods. Standard argument then implies that the relationship ends in one of the absorbing states with probability 1, and this shows that $\Pr(\tau < \infty) = 1$.

Part (i): when $p \geq p^{**}$, the PPE frontier is described as in part (i) of Lemma 5'. In this case, it is clear that the dynamics following the optimal equilibrium has only two steady states: permanent uncooperative delegation and termination. This proves part (i).

Next, suppose $p < p^{**}$. Notice that the relationship starts to the right of a, and if there exists t such that $\Pr(u_t < a) > 0$, then termination is a steady state with positive probability. Otherwise, termination is never reached and the only two steady states are permanent centralization and entrenchment. By Lemma 5' (part (ii) and (iii)), $u_l(\underline{u}_{CD}) \leq a$. It is then clear that $\Pr(u_t < a) = 0$ if and only if $u_l(\underline{u}_{CD}) = a$. We now show below that there exists p^* such that for all $p \leq p^*$, $u_l(\underline{u}_{CD}) = a$.

To do this we show that if $u_l(\underline{u}_{CD}, p') = a$ then $u_l(\underline{u}_{CD}, p'') = a$ for $p'' < p'$. Let s_0 be the slope between (a, a) and $(\underline{u}_{CD}, \pi(\underline{u}_{CD}))$, and s_1 be the slope between $(\bar{u}_{CD}, \pi(\bar{u}_{CD}))$ and (B, b). Note that both s_0 and s_1 depend on p. Now if $u_l(\underline{u}_{CD}, p') = a$, this implies that

$$s_0(p') \leq (1 - p')\pi'_-(a, p') + p'\pi'_-(u_h(\underline{u}_{CD}), p').$$

To show that $u_l(\underline{u}_{CD}, p'') = a$ for $p'' < p'$, it then suffices to show that

$$s_0\left(p''\right) \leq (1-p'')\pi'_-(a,p'') + p''\pi'_-(u_h(\underline{u}_{CD}), p'').$$

This is because we can use the standard argument to show that $\pi(u)$ is the unique fixed point of the operator

$$Tf\left(u\right) = \max\left\{\begin{array}{c}\max\left\{\pi_C\left(u\right), \pi_{D_C}\left(u\right), \pi_{D_U}\left(u\right), \pi_E\left(u\right)\right\}, \\ \max_{\alpha\in(0,1),u_1,u_2\in[0,B]}\left\{\alpha f\left(u_1\right) + \left(1-\alpha\right)f\left(u_2\right)\right\}\end{array}\right\},$$

which is monotone and nonexpansive (see an analgous argument in Step 2 in Lemma 5C).

A sufficient condition is then given by

$$\frac{\partial s_0\left(p\right)}{\partial p} > \partial\left((1-p)\pi'_-(a,p) + p\pi'_-(u_h(\underline{u}_{CD}), p)\right)/\partial p$$

for all $p < p'$. Notice that if $\partial s_0\left(p\right)/\partial p$ does not exist, we can replace it with the left derivative. To see that the left derivative exists, notice that $\pi(u)$ is weakly increasing in p for all u, so $s_0(p)$ is weakly increasing in all p, and therefore, the left derivative exists. Now

$$\begin{aligned} & \partial\left((1-p)\pi'_-(a,p) + p\pi'_-(u_h(\underline{u}_{CD}), p)\right)/\partial p \\ = \;& -\pi'_-(a,p) + (1-p)\frac{\partial\pi'_-(a,p)}{\partial p} + \pi'_-(u_h(\underline{u}_{CD}), p) + p\frac{\partial\pi'_-(u_h(\underline{u}_{CD}), p)}{\partial p} \\ \leq \;& p\frac{\partial\pi'_-(u_h(\underline{u}_{CD}), p)}{\partial p}, \end{aligned}$$

where the inequality follows because $\pi'_-(u_h(\underline{u}_{CD}), p) - \pi'_-(a,p) \leq 0$ by concavity of π and $\partial\pi'_-(a,p)/\partial p \leq 0$ because $\pi(a,p) = a$ for all $p < p^*$ and $\pi(u,p)$ is weakly increasing in p. It now follows that it suffices to show that

$$\frac{\partial s_0\left(p\right)}{\partial p} > p\frac{\partial\pi'_-(u_h(\underline{u}_{CD}), p)}{\partial p}.$$

To do this, notice that we can write

$$\pi'_-(u_h(\underline{u}_{CD}), p) = \alpha(p)s_0(p) + (1-\alpha(p))s_1(p)$$

for some $\alpha(p)$. Notice that $\partial s_1(p)/\partial p \leq 0$ since $\pi(u,p)$ increasing in p, which follows from an argument analgous to that in Step 2 in Lemma 5C. It follows that if $\partial\alpha(p)/\partial p < 0$, then

$$\begin{aligned} & p\frac{\partial\pi'_-(u_h(\underline{u}_{CD}), p)}{\partial p} \\ = \;& p\left(\frac{\partial\alpha(p)}{\partial p}s_0(p) + \alpha(p)\frac{\partial s_0(p)}{\partial p} + (1-\alpha(p))\frac{\partial s_1(p)}{\partial p} - \frac{\partial\alpha(p)}{\partial p}s_1(p)\right) \\ \leq \;& p\alpha(p)\frac{\partial s_0(p)}{\partial p} \\ \leq \;& \frac{\partial s_0(p)}{\partial p}. \end{aligned}$$

To show that $\partial\alpha(p)/\partial p < 0$, notice that $\alpha(p)$ is the probability that u_t falls into $[0, a]$ before $(\bar{u}_{CD}, B]$, where $\{u_t\}$ is a sequence of the agent payoffs starting at $u_1 = \underline{u}_{CD}$ and determined by the transition rule that

$$u_{t+1} = \begin{cases} u_h(u_t) & \text{with probability } p \\ u_l(u_t) & \text{with probability } 1-p \end{cases}.$$

Notice that we have the open bracket at u_1 and closed bracket at a because we are looking at the left derivatives here. In general, let $F(u,p)$ be the probability that u_t falls into $[0,a]$ first before $(\bar{u}_{CD}, B]$, and the sequence starts at u. In particular, let $F(u,p) = 1$ for $u \leq 0$ and $F(u,p) = 1 = 0$ for $u \geq B$. Now consider $p_1 < p_2$. For each p_i, $F(u,p_i)$ satisifes $F(u,p_i) = T_iF(u,p_i)$, where T_i is an operator on functions with ranges in $[a, B]$ satisfying

$$T_iF(u) = \begin{cases} \frac{u-a}{\underline{u}_{CD}-a}F(\underline{u}_{CD}) + \frac{\underline{u}_{CD}-u}{\underline{u}_{CD}-a} & 0 \leq u < \underline{u}_{CD} \\ p_iF(u_h(u)) + (1-p_i)F(u_\ell(u)) & \underline{u}_{CD} \leq u \leq \bar{u}_{CD} \\ \frac{B-u}{B-\bar{u}_{CD}}F(\bar{u}_{CD}) + \frac{u-\bar{u}_{CD}}{B-\bar{u}_{CD}}b & \bar{u}_{CD} < u \leq B \end{cases}$$

Notice that T_i is a monotone operator and has a unique fixed point for each p_i. Moreover, it is clear that $F(u,p_i)$ is decreasing in u for each p_i. Let $F(u,p_1)$ be the unique solution for $F(u,p_1) = T_1F(u,p_1)$. It is clear that $T_2F(u,p_1) \leq F(u,p_1)$ (since $F(u,p_1)$ is decreasing in u) for all u and it follows that $F(u,p_2) \leq F(u,p_1)$ for all u. In particular, $F(\underline{u}_{CD}, p_2) \leq F(\underline{u}_{CD}, p_1)$ and this proves that $\partial\alpha(p)/\partial p < 0$. This finishes showing that if $u_l(\underline{u}_{CD}, p') = a$ then $u_l(\underline{u}_{CD}, p'') = a$ for $p'' < p'$. As a result, there exists p^* such that for all $p \leq p^*$, $u_l(\underline{u}_{CD}) = a$ and there are two steady states in the long run.

Finally, we show that $p^* < p^{**}$. It is clear that $p^* \leq p^{**}$, so it suffices to rule out that $p^* = p^{**}$. Recall that at $p = p^{**}$, we have $\pi^T(a, p^{**}) = a$, and $\pi^T(u, p^{**}) = \pi(u, p^{**})$. We now show that there exists some small $\varepsilon > 0$ such that for $p = p^{**} - \varepsilon$, $u_l(\underline{u}_{CD}, p) < a$. Then by the argument above, we have $p^* < p^{**} - \varepsilon$. Define $u_a \equiv u_l^{-1}(a)$ and let $s(p^{**})$ be the slope between (a,a) and $(u_a, \pi(u_a, p^{**}))$. Now suppose to the contrary that that for all $\varepsilon > 0, u_l(\underline{u}_{CD}, p^{**} - \varepsilon) = a$. Let $s(p)$ be the slope between (a,a) and $(u_a, \pi(u_a, p))$. Notice that for all $p < p^{**}$, $\pi(a,p) = a$. In addition, $\pi(u_a, p) \leq \pi(u_a, p^{**})$ because π is weakly increasing in p. This implies that $s(p) \leq s(p^{**})$ for all $p < p^{**}$. Now there are two cases to consider. First, $\pi(u, p^{**})$ is not a straight line between a and u_a, i.e., there exists a $u \in (a, u_a)$ such that $\pi(u, p^{**}) > a + s(p^{**})(u-a)$. In this case, we have

$$\pi(u, p^{**}) > a + s(p^{**})(u-a) \geq \lim_{\varepsilon \to 0} \pi(u, p^{**} - \varepsilon),$$

which violates the continuty of $\pi(u,p)$ in p. This is a contradiction.

In the second case, $\pi(u, p^{**})$ is a straight line between a and u_a. Notice that since $u_l(u_a) = a$, we have $a < u_a \leq \bar{u}_{CD}$. Lemma 5' then implies that for $u \in (a, u_a)$,

$$\pi'(u, p^{**}) = p\pi'(u_h(u), p^{**}) + (1-p)\pi'(u_l(u), p^{**}).$$

Since $\pi'(u, p^{*})$ is a constant for $u \in (a, u_a)$ and because π is concave, we have that π' is again a constant for all $u \in (u_l(a), a)$. Notice that

$$a - u_l(a) = (u_a - a)/\delta > u_a - a.$$

Now either $u_l(a) \geq \underline{u}_{CD}^A \equiv (1-\delta) B$ or $u_l(a) < \underline{u}_{CD}^A$. Recall that π' is a constant for $u \leq \underline{u}_{CD}^A$, and denote the slope of π in this region as s_0. Notice that if $u_l(a) < \underline{u}_{CD}^A$, we must have $u_l(a) = 0$ since $\pi'_{-}(u) < s_0$ for all $u > \underline{u}_{CD}^A$. Now if $u_l(a) \geq \underline{u}_{CD}^A$, we then have that π' is a constant for $u \in (u_{ll}(a), u_l(a))$, and repeating the same argument, we must have that $u_{l^n}(a) = 0$ for some n.

Next, we claim that $u_h(u_1) \geq \bar{u}_{CD} \equiv (1-\delta) b + \delta B$. Notice that π' is a constant for $u \in (u_h(\underline{u}_{CD}^A), u_h(\bar{u}_{CD}^A))$, and

$$u_h(\bar{u}_{CD}^A) - u_h(\underline{u}_{CD}^A) = \left(\bar{u}_{CD}^A - \underline{u}_{CD}^A\right)/\delta.$$

If $u_h(u_1) \leq \bar{u}_{CD}$, we must then also have $u_h(\bar{u}_{CD}^A) \leq \bar{u}_{CD}$ (because $\pi'_{-}(u)$ for $u < \bar{u}_{CD}$ is strictly larger than the slope of π for $u \geq \bar{u}_{CD}$.) But if π' is a constant for $u \in (u_h(u_1), u_h(\bar{u}_{CD}^A))$, same argument as above implies that π' is a constant for $u \in (u_{lh}(\underline{u}_{CD}^A), u_{lh}(\bar{u}_{CD}^A))$, which is a line segment with length $\left(\bar{u}_{CD}^A - \underline{u}_{CD}^A\right)/\delta^2$. Applying the same argument as before, this would imply that the distance between $\bar{u}_{CD}^A$ and$\underline{u}_{CD}^A$ is bigger than $\left(\bar{u}_{CD}^A - \underline{u}_{CD}^A\right)/\delta^2$, which is a contradiction.

Finally, since $\underline{u}_{CD}^A = (1-\delta) B$, $u_h(\underline{u}_{CD}^A) \geq \bar{u}_{CD}$ then implies that

$$\frac{(1-\delta)(B-b)}{\delta} \geq B - (1-\delta)b > B - a,$$

which contradicts Assumption X ($\frac{(1-\delta)(B-b)}{\delta} \leq B - a$.)$\blacksquare$