Repeated Delegation

by *Elliot Lipnowski & João Ramos*[*]

NEW YORK UNIVERSITY

November, 2015

Abstract

We study an ongoing relationship of delegated decision making. Facing a stream of projects to potentially finance, a principal must rely on an agent to assess the returns of different opportunities; the agent has lower standards, wishing to adopt every project. In equilibrium, the principal allows bad projects in the future to reward fiscal restraint by the agent today. We fully characterize the equilibrium payoff set, showing that Pareto optimal equilibria can be implemented via a two-regime 'Dynamic Capital Budget'. We show that, rather than backloaded rewards—a prevalent feature of dynamic agency models with greater commitment power—our Pareto optimal equilibria feature an inevitable loss of autonomy for the agent as time progresses. This drift toward conservatism speaks to the life cycle of an organization: as it matures, it grows less flexible but more efficient.

JEL codes: C73, D23, D73, D82, D86, G31

[*]Email: `elliot.lipnowski@nyu.edu` and `ramos@nyu.edu`

We thank Joyee Deb, David Pearce, and Ennio Stacchetti for valuable insights and constant encouragement. This work has also benefited from discussions with Heski Bar-Isaac, Dhruva Bhaskar, V. Bhaskar, Adam Brandenburger, Sylvain Chassang, Brendan Daley, Peter DeMarzo, Ross Doppelt, Ignacio Esponda, Eduardo Faingold, Johannes Hörner, R. Vijay Krishna, Vijay Krishna, Laurent Mathevet, Debraj Ray, Alejandro Rivera, Ariel Rubinstein, Evan Sadler, Tomasz Sadzik, Maher Said, Alex Wolitzky, Chris Woolnough, Sevgi Yuksel, and seminar participants at New York University, EconCon, University of Waterloo, Midwest Economic Theory, the Canadian Economic Theory Conference, and the SITE dynamic games workshop. The usual disclaimer applies.

1 Introduction

Many economic activities are arranged via delegated decision making. In practice, those with the necessary information to make a decision may differ—and, indeed, have different interests—from those with the legal authority to act. Such relationships are often ongoing, consisting of many distinct decisions to be made over time, with the conflict of interest persisting throughout. A state government that funds local infrastructure may be more selective than the local government equipped to evaluate its potential benefits. A university bears the cost of a hired professor, relying on the department to determine candidates' quality. The Department of Defensefunds specialized equipment for each of its units, but must rely on those on the ground to assess their need for it. Our focus is on how such repeated delegation should optimally be organized, and on how the relationship evolves over time.

Beyond the limits of monetary incentives, formal contingent contracting may be difficult for two reasons. First, it may be impractical for the informed party to produce verifiable evidence supporting its recommendations. Second, it might be unrealistic for the controlling party to credibly cede authority in the long-run. Even so, the prospect of a future relationship may align the actors' interests: both parties may be flexible concerning their immediate goals, with a view to a healthy relationship.

We study an infinitely repeated game in which a principal ("she") with full authority over a decision to be made in an uncertain world; she relies on an agent ("he") to assess the state. Each period, the principal must choose whether or not to initiate a project, which may be good (i.e. high enough value to offset its cost) or bad. The principal herself is ignorant of the current project's quality, but the agent can assess it. The players have partially aligned preferences: both prefer a good project to any other outcome, but they disagree on which projects are worth taking. The principal wishes to fund only good projects, while the agent always prefers to invest in any project. For instance, consider the ongoing relationship between local and state governments. Each year, a county can request state government funds for the construction of a park. The state, taking into account past funding decisions, decides whether or not to fund it. The park would surely benefit the county, but the state must weigh this benefit against the money's opportunity cost. To assess this tradeoff, the state relies on the county's local expertise. We focus on the case in which the principal needs the agent: the ex-ante expected value of a project is not enough to offset its cost. Thus, if the county were never selective in its proposals, the state would never want to fund the park.

To delegate—to cede control at the ex-ante stage—entails some vulnerability. If our principal wants to make use of the agent's expertise, she must give him the liberty to act. In funding a park, the state government risks wasting taxpayer money. Acting on a county's recommendation, the state won't know whether the park is truly valuable to the community, even after it is built. Furthermore, if the state makes a policy of funding each and every park the county requests, then it risks wasting a lot of money on many unneeded parks. This vulnerability limits the freedom that the agent can expect from the principal in the future. The state government cannot credibly reward a county's fiscal restraint today by promising carte blanche in the future.

The present conflict of interest would be resolved if the principal could sell permanent control to the agent.[1] In keeping with our leading applications, we focus on the repeated interaction without monetary transactions.[2] The Department of Defense, for example, is unlikely to ask soldiers to pay for their own body armor.

Our first main result says that the delegation relationship drifts toward conservatism as time progresses. In the early stages of the relationship, the principal always cedes control to the agent, who in turn adopts every good project that arrives but occasionally adopts bad projects as well. At this stage, the agent enjoys a high per-period value, but the principal faces a low yield, in that the average adopted project has a low value per unit cost. Eventually, the agent's goodwill permanently runs dry. After this happens, the principal rarely delegates; the agent may adopt some good projects if allowed, but will never again squander an opportunity on a bad project.

Thus, combining three sensible ingredients—limited information, limited lia-

[1]This standard solution is sometimes called "selling the firm." For instance, see p. 482 of **?**.

[2]This assumption is stronger than needed. As long as the agent enjoys limited liability, our main results are qualitatively unchanged.

bility, and limited commitment—yields surprising conclusions for the evolution of the relationship. Indeed, if the agent's cost of exercising restraint were observable, we would expect him to be rewarded over time, as in **?**. If, maintaining unobservability, the principal could commit to future choices, the agent would face divergent long-run outcomes, as in **?**. Our qualitatively different results yield testable implications for the life cycle of an institution: *as an organization matures, it grows less flexible but more efficient*.

Our second main result is a complete characterization of the equilibrium payoff set. Given the lack of a general method to explicitly compute this set for a repeated game of incomplete information (at fixed discounting), we must combine the recursive methods of **?** with several specific features of our stage game to arrive at a tractable problem. For instance, because the principal's actions are immediately perfectly observed, her incentive constraint amounts to a lower bound on her continuation payoff. This key step allows us to represent an important piece of our value set via dynamic programming, enabling an analytical solution.

Using the above, we obtain our third main result, which uncovers the form of the optimal intertemporal delegation rule. In keeping with the basic intuition (as in, say, **?**) that linking decisions aligns incentives, delegation is organized via a budgeting rule. The uniquely[3] principal-optimal equilibrium, the **Dynamic Capital Budget**, comprises two distinct regimes. At any time, the parties engage in either Capped Budgeting or Controlled Budgeting.

In the **Capped Budget** regime, the principal always delegates, and the agent initiates all good projects that arrive. At the relationship's outset, the agent has an expense account for projects, indexed by an initial balance and an account balance cap. The balance captures the number of projects that the agent could adopt immediately without consulting the principal. Any time the agent takes a project, his balance declines by 1. While the agent has any funds in his account, the account accrues interest. If the agent takes few enough projects, the account will grow to its cap. At this balance, the agent is still allowed to take projects, but his account grows no larger (even if he waits). Not being rewarded for fiscal restraint, the agent

[3] More precisely, the two-regime structure and the exact character of the Capped Budget regime are uniquely required by optimality.

immediately initiates a project, and his balance again declines by 1.

If the agent overspends, a **Controlled Budget** regime begins: the principal only delegates sometimes in this regime, but the agent never adopts a bad project again. This inevitable outcome—indeed Theorem 1 tells us Capped Budgeting can only be temporary—is a unfortunate one, Pareto inferior to the principal-optimal equilibrium. This dismal future, with the threat it entails, is necessary to sustain productive delegation in the earlier stages of the relationship.

Related Literature This paper studies delegated decision making, a topic initiated by **?**,[4] focusing on the tradeoff a principal faces between leveraging an agent's private information and shielding herself from his conflicting interests. A key insight from this literature[5] is that some ex-post inefficiency may be optimal, providing better incentives to the agent ex-ante.

We join a growing literature on dynamic delegation. Most closely related[6] is the contemporaneous work by **?**, studying optimal dynamic mechanisms without money in a world of partially persistent valuations, in which the principal has commitment power. The optimal mechanism in their model generates path dependence in long-run outcomes: in contrast to our model, the agent may receive his first-best outcome indefinitely.

Our model speaks to the relational contracting literature,[7] which focuses on relationships in which formal contracting is limited, and all incentives—and the credibility of promises that provide those incentives—are anchored to the future value of the relationship. In contemporaneous work, **?** focus on a repeated trust game (the principal deciding whether to trust the agent) preceded at every stage by a simultaneous entry decision. The entry decision (together with a lower bound on patience which they invoke) gives their principal effective commitment power,[8]

[4] Also, see **?** and the thorough review therein.

[5] See **?** and **?**, for example.

[6] Also related is **?**, characterizing the optimal contract for a principal who delegates investment choices and has a costly state verification technology, monetary transfers, and commitment power. See **?** and **?** as well.

[7] See **?**, **?**, and **?**.

[8] Indeed, one can show that the optimal contracting analogue of their leading model—wherein the principal faces no incentive-compatibility constraints—generates the same path of play as their

generating history dependence similar to that in **?**: random early outcomes have long-lasting consequences.[9] In our model, irrespective of how patient the players are, the agent cannot credibly be rewarded in the long-run.

Our results add to the literature on relationship building under private information. While focusing on a different misalignment of preferences, **?** and **?** look at a model of trading favors,[10] in which a player's opportunity to do a favor is private. The form of the conflict of interests is qualitatively different there: unlike in our model (wherein adopting good projects benefits everybody), every action that benefits one player harms another in the stage game. **?** show that the relationship can benefit from varying incentives based on both action and inaction, a feature which reappears in our model. Their simulations indicate that forgiveness is a feature of every Pareto optimal equilibrium. In our model of partially aligned preferences, there are fundamental limits to forgiveness: our players eventually permanently face a Pareto dominated outcome.

The use of allocation of future responsibility as an incentivizing device is familiar to the dynamic corporate finance literature. In **?**, for instance,[11] the principal commits to investment choices and monetary transfers to the agent, who privately acts to reduce the chance of large losses for the firm. While our setting is considerably different, their optimal contract and our Pareto optimal equilibria exhibit similar dynamics: our "funny money" balance plays the same role in our model that real sunk investment plays in theirs.

Lastly, there is a connection between the present work and the literature on linked decisions. **?**[12] show that, given a large number of physically independent decisions, the ability to connect them across time helps align incentives. **?** shows that a principal with commitment power optimally commits to a budgetary rule to discipline an agent with state-independent preferences. In our model—without such commitment power, and with partial alignment of preferences—dynamic budgeting

principal-optimal equilibrium.

[9] See **?** as well.

[10] Also see **?**, **?**, **?**, and **?** on the same theme. Our model features one-sided incomplete information like **?**, but two-sided lack of commitment like the rest.

[11] Also see **?** and **?**.

[12] See **?** and **?** too.

remains optimal but is tempered by the principal's need for credibility.

Structure of the paper The remainder of the paper is structured as follows. Section 2 presents the model and introduces a convenient language for discussing players' incentives in our model. In Section 3, we discuss aligned equilibria, i.e. those in which no bad projects are adopted; we characterize the class and show that such equilibria are necessarily inefficient. Section 4 studies the dynamics of the relationship, showing that it becomes more conservative as time progresses. The heart of the paper is Section 5, wherein we fully characterize the equilibrium payoff set and the path of play of efficient equilibria. In Section 6, we present the Dynamic Capital Budget, a straightforward mechanism that implements Pareto efficient equilibria. In Section 7, we discuss some possible extensions of our model. Final remarks follow in Section 8.

2 The Model

We consider an infinite-horizon two-player ($\mathcal{P}$rincipal and $\mathcal{A}$gent) game, played in discrete time. Each period, the principal chooses whether or not to delegate a project adoption choice to the agent. Conditional on delegation, the agent privately observes which type of project is available and then publicly decides whether or not to adopt it. At the time of its adoption, a project of type θ generates an agent payoff of θ. Each project entails an implementation cost of c, to be borne solely by the principal; thus, a project yields a net utility of $\theta - c$ to the principal. Notice that the cost is independent of the project's type. In particular, the difference between the agent's payoffs and the principal's payoffs doesn't depend on the agent's private information. We interpret this payoff structure as the principal innately caring about the agent's (unobservable) payoff, in addition to the cost that she alone bears. While the university's president cannot expertly assess a specialized candidate, she still wants the physics department to hire good physicists. The state government can't assess the added value of each local public project, but it still values the benefit that a project brings to the community. The principal cares about the value generated by a project, although she never observes it.

While the players rank projects in the same way, the key tension in our model is a disagreement over which projects are worth taking. The agent cares only about the benefit generated by a project, while the principal cares about said benefit net of cost; we find revenue and profit to be useful interpretations of the players' payoffs.

$\mathcal{P}$ and $\mathcal{A}$ share a common discount factor $\delta \in (0,1)$, maximizing *expected discounted profit* and *expected discounted revenue*, respectively. So, if the available project in each period $k \in \mathbb{Z}_+$ is θ_k and projects are adopted in periods $\mathcal{K} \subseteq \mathbb{Z}_+$, then the principal and agent get profit and revenue,

$$\Pi = (1-\delta)\sum_{k\in\mathcal{K}} \delta^k(\theta_k - c) \text{ and } V = (1-\delta)\sum_{k\in\mathcal{K}} \delta^k \theta_k, \text{ respectively.}$$

Each period, $\mathcal{P}$ and $\mathcal{A}$ play the following stage game:

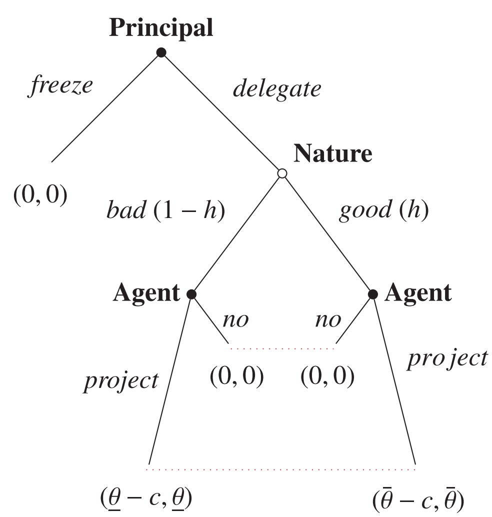

Figure 1: The principal observes the agent's choices but not project quality.

First, $\mathcal{P}$ publicly decides whether to *freeze* project adoption or to *delegate* it. If $\mathcal{P}$ freezes, then no project is adopted and both players accrue no payoffs. If $\mathcal{P}$ delegates, then $\mathcal{A}$ privately observes which type of project is available and decides whether or not to initiate the available project. The current period's project is good (i.e. of type $\bar{\theta}$) with probability $h \in (0,1)$ and bad (i.e. of type $\underline{\theta}$) with complementary probability. If the agent initiates a project of type θ, payoffs $(\theta - c, \theta)$ accrue to the players. The principal observes whether or not the agent initiated a project, but she never sees the project's type.

Notation. *Let* $\theta_E := (1-h)\underline{\theta} + h\bar{\theta}$ *be the* ***ex-ante expected project value***.

Throughout the paper, we maintain the following assumption:

Assumption 1.

$$0 < \underline{\theta} < \theta_E < c < \bar{\theta}.$$

Assumption 1 characterizes the preference misalignment between agent and principal. Since $\underline{\theta} - c < 0 < \bar{\theta} - c$, the principal prefers good projects to nothing, but prefers inactivity to bad projects. Given $0 < \underline{\theta} < \bar{\theta}$, the agent prefers any project to no project, but also prefers good ones to bad ones. The condition $\theta_E < c$ (interpreted as an assumption that good projects are scarce) says that the latter effect dominates, and the conflict of interest prevails even ex-ante: the principal prefers a freeze to the average project. A good enough physicist is rare; the university finds hiring worthwhile only if it can rely on the department to separate the wheat from the chaff. If the players interacted only once, the department would not be selective. Accordingly, the stage game has a unique sequential equilbrium: the principal freezes, and the agent takes a project if allowed.

Before proceeding to characterize the equilibrium payoff set, we document one more piece of notation, and we make an interpretable assumption which will guarantee existence of interesting equilibria.

Notation. *Let* $\omega := h(\bar{\theta} - \underline{\theta}) = \theta_E - \underline{\theta}$ *be the* ***marginal value of search***.

The constant ω captures the marginal option value of seeking another project rather than taking an existing bad project.

Assumption 2.

$$\delta\omega \geq (1-\delta)\underline{\theta} \textit{ or, equivalently, } \omega \geq (1-\delta)\theta_E.$$

Henceforth, we take Assumption 2 as given. This assumption, which can equivalently be expressed as a lower bound on the discount factor δ, essentially says that intertemporal tradeoffs alone are enough to incentivize good behavior from the agent. If the agent is sufficiently patient, the marginal value of searching for a good project outweighs the myopic benefit of an immediate bad project.

While Assumption 1 (in particular, $c > \theta_E$) tells us that the only *stationary* equilibrium of our repeated game is an unproductive one—i.e. never has any (good or bad) projects adopted—Assumption 2 guarantees existence of *some* productive equilibria. Indeed, it implies the following is an equilibrium: before the first project is adopted, the principal delegates and the agent adopts only good projects; after the first project is adopted, they play the stage game equilibrium. Simple algebra shows that Assumption 2 exactly delivers the agent's incentive constraint to willingly resist a bad project.

Lastly, we make the following assumption, another lower bound on the discount factor, for convenience.

Assumption 3.

$$(1-\delta)(c-\theta_E) \leq \delta\left(1-\frac{c}{\bar{\theta}}\right)\omega.$$

Assumption 3, whose sole purpose is to guarantee existence of some equilibrium with certain initial project adoption,[13] makes our main results less cumbersome to state; the vast majority of our work proceeds readily without it.

2.1 Equilibrium Values

Throughout the paper, *equilibrium* will be taken to mean perfect semi-public equilibrium (PPE), in which the players respond only to the public history of actions and (for the agent) current project availability. This definition in similar in spirit to that in **?** and in **?**.

Definition 1. *Each period, one of three public outcomes occurs: the principal freezes; the principal delegates and the agent initiates no project; or the principal delegates and the agent initiates a project. A period-k* ***public history****, h_k, is a sequence of k public outcomes (along with realizations of public signals). The agent*

[13] Indeed, we could replace Assumption 3 with the endogenous assumption, "There exists an equilibrium in which a project is certain to be adopted at time zero," and leave all our results unchanged.

has more relevant information when making a decision. A time-k ***agent semi-public history*** *is* $h_k^{\mathcal{A}} = (h_k, D, \theta_k)$*, where* h_k *is a public history, D is a principal decision to delegate, and* θ_k *is a current project type.*

A ***principal public strategy*** *specifies, for each public history, a probability of delegation. An* ***agent semi-public strategy*** *specifies, for each agent semi-public history, a probability of project adoption.*

A ***perfect (semi-)public equilibrium*** *(PPE) is a sequential equilibrium in which the principal plays a public strategy, while the agent plays a semi-public strategy.*

Every equilibrium entails an expected discounted number of adopted good projects $g = (1-\delta)\mathbb{E}\sum_{k\in\mathcal{K}} \delta^k \mathbf{1}_{\{\theta_k=\bar{\theta}\}}$ and an expected discounted number of adopted bad projects $b = (1-\delta)\mathbb{E}\sum_{k\in\mathcal{K}} \delta^k \mathbf{1}_{\{\theta_k=\underline{\theta}\}}$, where $\mathcal{K} \subseteq \mathbb{Z}_+$ is the realized set of periods in which the principal delegates and the agent adopts a project. Given those, one can compute the agent value/**revenue** as

$$v = \bar{\theta}g + \underline{\theta}b$$

and the principal value/**profit** as

$$\pi = (\bar{\theta} - c)g - (c - \underline{\theta})b.$$

For ease of bookkeeping, it is convenient to track equilibrium-supported revenue v and bad projects b, both in expected discounted terms. The vector (v, b) encodes both agent value v and principal profit

$$\begin{aligned}
\pi(v,b) &:= (\bar{\theta} - c)g - (c - \underline{\theta})b \\
&= (\bar{\theta} - c)\frac{v - \underline{\theta}b}{\bar{\theta}} - (c - \underline{\theta})b \\
&= \left(1 - \frac{c}{\bar{\theta}}\right)v - c\left(1 - \frac{\underline{\theta}}{\bar{\theta}}\right)b.
\end{aligned}$$

In (v, b) space, the principal's indifference curves are of slope $\xi := \frac{\bar{\theta}-c}{c(\bar{\theta}-\underline{\theta})}$.

Toward a Characterization The main objective of this paper is to characterize the set of equilibrium-supported payoffs,

$$\mathcal{E}^* := \{(v, b) : \exists \text{ equilibrium with revenue } v \text{ and bad projects } b\} \subseteq \mathbb{R}^2_+.$$

Throughout the paper, we make extensive use of two simple observations about our model. First, notice that $(0, 0) \in \mathcal{E}^*$, since the profile σ^{static}, in which the principal always freezes and the agent takes every permitted project, is an equilibrium. Said differently, there is always an unproductive equilibrium—i.e. one with no projects. That this equilibrium provides min-max payoffs makes our characterization easier. Second, as the following lemma clarifies, off-path strategy specification is unnecessary in our model. If appropriate on-path incentive constraints are satisfied, we can alter behavior off-path to yield an equilibrium. With the lemma in hand, we rarely specify off-path behavior in a given strategy profile.

Lemma 1. *Fix a strategy profile σ, and suppose that:*

1. *The agent has no profitable deviation from any on-path history.*
2. *At all on-path histories, the principal has nonnegative continuation profit.*

Then, there is an equilibrium $\tilde{\sigma}$ that generates the same on-path behavior (and, therefore, the same value profile).

Proof. Let σ^{static} be the stage Nash profile—i.e. the principal always freezes, and the agent takes a project immediately whenever permitted. Define $\tilde{\sigma}$ as follows:

- On-path (i.e. if $\mathcal{P}$ has never deviated from σ), play according to σ.
- Off-path (i.e. if $\mathcal{P}$ has ever deviated from σ), play according to σ^{static}.

The new profile is incentive-compatible for the agent: off-path because σ^{static} is, on-path because σ is. It is also incentive-compatible for the principal: off-path because σ^{static} is, on-path because σ is and has nonnegative continuation profits while σ^{static} yields zero profit. □

Dynamic Incentives If we aim to understand the players' incentives at any given moment, we must first understand how their future payoffs respond to their current choices. To describe the law of motion of revenue v [or, respectively, bad projects b], we keep track of:

- $\check{v}$, next period's agent continuation value if the principal freezes today;
- v', the agent continuation value if the principal delegates today and the agent abstains from project adoption;
- $\tilde{v}$, the agent continuation value if a project is undertaken today; and
- $\check{b}, b', \tilde{b}$ analogously for continuation (expected discounted) bad projects.

Observe that the continuation values cannot depend on the quality of the adopted project (nor can the laws of motion depend on availability of forgone projects), which is not publicly observable. Finally, describe the players' present actions (after any public randomization occurs) as follows:

- The principal makes a delegation choice $p \in [0,1]$, the probability with which she delegates to the agent in the current period.
- The agent chooses $\bar{a} \in [0,1]$ and $a \in [0,1]$, the probabilities which he currently initiates available good and bad projects, respectively, conditional on being allowed to.

Appealing to self-generation arguments, as in **?**, and to Lemma 1, equilibrium is characterized by the following three conditions:

1. Promise keeping:

$$\begin{aligned}(v,b) \quad = \quad & (1-p)\delta(\check{v},\check{b}) + ph\left\{\bar{a}\left[(1-\delta)(\bar{\theta},0)+\delta(\tilde{v},\tilde{b})\right]+(1-\bar{a})\delta(v',b')\right\} \\ & +p(1-h)\left\{a\left[(1-\delta)(\underline{\theta},1)+\delta(\tilde{v},\tilde{b})\right]+(1-a)\delta(v',b')\right\}\end{aligned}$$

   We decompose continuation outcomes (v,b) after any period into what happens in each of four events—the agent finds and invests in a good project; the agent finds and invests in a bad project; no project is adopted by the agent's choice; no project is adopted on the principal's authority—weighted by their equilibrium probabilities.

2. Agent incentive compatibility:

$$\begin{aligned} \delta[v' - \tilde{v}] &\geq (1-\delta)\underline{\theta} \ \text{ if } a < 1, \\ \delta[v' - \tilde{v}] &\leq (1-\delta)\bar{\theta} \ \text{ if } \bar{a} > 0, \\ \bar{a} &\geq a. \end{aligned}$$

   If the agent is willing to resist taking a project immediately ($a < 1$), it must be that the punishment $v - \tilde{v}$ for taking a project is severe enough to deter the $\underline{\theta}$ myopic gain; similarly, if the agent is to take some good projects ($\bar{a} > 0$), the same punishment $v - \tilde{v}$ cannot be too draconian. Lastly, no future play can make bad projects more appealing to the agent than good projects.

3. Principal participation:[14]

$$\pi(v, b) \geq 0.$$

   The principal could, at any moment, unilaterally move to a permanent freeze and secure herself a profit of zero. Therefore, at any history, she must be securing at least that much in equilibrium.

3 Aligned Equilibrium

We have established that our game has no productive stationary equilibrium. If the principal allows history-independent project adoption, the agent cannot be stopped from taking limitless bad projects. In the present section, we ask whether this core tension can be resolved by allowing non-stationary equilibria.

Definition. *An **aligned equilibrium** is an equilibrium in which no bad projects are ever adopted on-path.*

[14] This is, of course, a relaxation of the principal's incentive constraint. Even in light of Lemma 1, the principal must be indifferent between freezing and delegating for $p \in (0, 1)$ to be incentive-compatible. In the appendix, we show that it is without loss for the principal to never privately mix. With this in mind, the present relaxed incentive constraint is all that matters.

As an example, consider the following straightforward strategy profile σ_ϕ, described by a firing probability $\phi \in [0, 1]$. The principal initially delegates, and the agent adopts the first project if and only if it is of type $\bar{\theta}$. If the agent abstains today, σ_ϕ begins again tomorrow. If the agent takes a project today, then tomorrow a public coin is tossed: the continuation play is either σ_ϕ (with probability $1 - \phi$) or σ^{static} (with probability ϕ).

So the principal delegates to the agent only temporarily, and the agent always serves the principal's interests. If and only if the risk ϕ of firing is high enough to make an unlucky (i.e. facing a $\underline{\theta}$ project) agent exercise restraint—which Assumption 2 guarantees will happen for high enough ϕ—this profile will be an equilibrium. The optimal equilibrium in this class then leaves the incentive constraint $\phi\delta v_\phi \geq (1-\delta)\underline{\theta}$ binding, as lowering ϕ yields a Pareto improvement. For this choice of firing probability,

$$v_\phi = (1 - \delta)h\bar{\theta} + \delta(1 - \phi h)v_\phi = \delta v_\phi + h(1 - \delta)(\bar{\theta} - \underline{\theta}) \implies v_\phi = \omega.$$

This simple class illuminates the forces at play in our model. The principal wants good projects to be initiated, but she cannot afford to give the agent free rein. If she wants to stop him from investing in bad projects, she must threaten him with mutual surplus destruction. Subject to wielding a large enough stick to encourage good behavior, she efficiently wastes as little opportunity as possible. The university does not want to deprive the physics department of needed faculty, and so it should limit them only enough to discipline fiscal restraint.

Remarkably, this sensible equilibrium is in fact optimal among all aligned equilibria.

Proposition 1 (Aligned Optimum)**.** *The highest agent value in any aligned equilibrium*[15] *is ω, the marginal value of search.*

The intuition behind the theorem, formally proven via Lemma 9 in the Appendix, is straightforward. If the agent's current expected value exceeds his marginal value of search, then asking an agent to search (i.e. to exercise restraint when facing

[15] Absent Assumption 2, there can be no aligned equilibrium yielding $v > 0$.

a bad project) requires growth in his continuation value. Eventually, after enough bad luck, this promised value can only be fulfilled by allowing some bad projects.

Corollary 1. *The aligned optimal profit is* $\omega\left(1 - \frac{c}{\theta}\right)$*, which is independent of players' patience and strictly below the principal's first-best profit* $h(\bar{\theta} - c)$*.*

To keep the agent honest, the principal must destroy some future surplus when the agent overspends. As the players become more patient, the punishment must be made more draconian to provide the same incentives. When only good projects are adopted, this added severity punishes the principal too, so that no welfare gains are achieved.

We know there is an equilibrium providing the principal with payoffs near her first-best—and thus Pareto dominating all aligned equilibria, by Corollary 1—if the players are sufficiently patient.[16] The Corollary therefore tells us that aligned equilibria are inadequate: bad projects are a necessary ingredient of a well-designed relationship.

4 Drifting toward Conservatism

The previous section showed us that a relationship of ongoing delegation should, at least for some parameter values, admit some bad projects. This considerably complicates our analysis for at least two reasons. First, there is an extra degree of freedom in providing agent value: not just when to have the agent take projects, but also which types to have him adopt. Second, we must now worry about the principal's incentives: too many bad projects would violate her participation constraint.

Our ultimate goal, which we will achieve in Section 5, is to characterize Pareto efficient equilibria and the value they produce. In the present section, we derive some key qualitative properties of the path of play of such equilibria.

The next theorem describes the dynamics of an efficient relationship between our two players. In the early stages of the relationship, no good project is missed, but some bad ones are adopted: the per-period revenue is high, but the average

[16]This fact can be deduced from **?**, by first removing a dominated stage game strategy (the agent only adopting bad projects) and then verifying indentifiability.

project taken in this stage has a return on investment $< \frac{\bar{\theta}}{c}$. As the relationship progresses and random project types unfold, the regime eventually changes. In the later stages of the relationship, some good projects are missed, but no bad projects are adopted: the per-period revenue is now lower, but the average project has a higher return on investment. Said differently, the players transition from a high-revenue, low-yield existence to a low-revenue, high-yield existence.

Theorem 1 (Relationship Dynamics)**.** *In any Pareto efficient equilibrium, there is (with probability 1) a finite time K such that:*

- *In period $k < K$: the agent's continuation value is $> \omega$, the principal delegates, and the agent adopts the current project if it is good.*
- *In period $k \geq K$: The agent's continuation value is $\leq \omega$, and the agent does not adopt the current project if it is bad.*

Proof. The idea is that, with high enough probability, the agent's value falls over time. We prove the theorem through a series of smaller claims. First note that $\bar{v} < \theta_E$. Indeed, the unique b such that (θ_E, b) is feasible is $b = 1 - h$, yielding negative profit to the principal. Now, let $\epsilon := \frac{1-\delta}{\delta}\left(\frac{\theta_E - \bar{v}}{2}\right) > 0$ and $q := \frac{\epsilon}{\bar{v}+\epsilon} \in (0, 1)$.

Claim 1: If the agent's continuation value is v before a public randomization, then the value $\hat{v}$ after public randomization is $\leq v + \epsilon$ with probability $\geq q$.
Proof: Letting $\hat{q} := \mathbb{P}\{\hat{v} \leq v + \epsilon\}$ yields $v = (1 - \hat{q})\mathbb{E}\{\hat{v}|\hat{v} > v + \epsilon\} + 1 - \hat{q}\mathbb{E}\{\hat{v}|\hat{v} \leq v + \epsilon\} \geq (1 - \hat{q})(v + \epsilon) + \hat{q}0 = (v + \epsilon) - \hat{q}(v + \epsilon)$, so that $\hat{q} \geq \frac{\epsilon}{v+\epsilon} \geq q$.

Claim 2: If (i) $(v, b), (v^1, b^1)$ are equilibrium payoffs with $(v, b) \in \text{co}\{(v^1, b^1), (0, 0)\}$; (ii) (v, b) is some on-path continuation value pair for some Pareto-efficient equilibrium; and (iii) $v \geq \omega$; then $(v, b) = (v^1, b^1)$.
Proof: If not, then $(v, b) = (\lambda v^1, \lambda b^1)$ for some $\lambda \in (0, 1)$. There is a unique $(v', b') \in \text{co}\{(v^1, b^1), (\omega, 0)\}$ with $v' = v$. By construction, $b' < b$, so that the equilibrium—$\mathcal{E}^*$ being convex—value pair (v', b') provides the exact same agent value as (v, b) but at a strictly higher principal value. Replacing the continuation play with some continuation equilibrium providing (v', b') maintains equilibrium, contradicting Pareto optimality of the original equilibrium.

Claim 3: If the agent's continuation value is $v \geq \omega$ after public randomization on the path of some Pareto efficient equilibrium, then the tomorrow's continuation value following a good project today is certain to be $\leq v - 2\epsilon$.

Proof: Suppose the agent's continuation value is $v \geq \omega$ after public randomization; let b be the expected discounted number of bad projects. Let stage play $p, \bar{a}, a$, and continuation payoffs $\check{v}, v', \tilde{v}, \check{b}, b', \tilde{b}$ be as described in the previous section.

First, notice that $(v, b) = (1-p)(1-\delta)(0,0) + [1-(1-p)(1-\delta)](v^1, b^1)$ for some $(v^1, b^1) \in \mathcal{E}$. Therefore, $p = 1$ by Claim 2; the principal must be delegating.

Now, if $\bar{a} < 1$ then agent IC would tell us $a = 0$, so that

$$\begin{aligned}
(v,b) &= h\left\{\bar{a}\left[(1-\delta)(\bar{\theta},0) + \delta(\tilde{v},\tilde{b})\right] + (1-\bar{a})\delta(v',b')\right\} + (1-h)\delta(v',b') \\
&= (1-\bar{a})\delta(v',b') + \bar{a}\left\{h\left[(1-\delta)(\bar{\theta},0) + \delta(\tilde{v},\tilde{b})\right] + (1-h)\delta(v',b')\right\} \\
&\in (1-\bar{a})\delta(v',b') + \bar{a}\mathcal{E}^*.
\end{aligned}$$

Again by Claim 2, this cannot be. So $\bar{a} = 1$; the agent takes every good project. Hence,

$$\begin{aligned}
v &= h\left[(1-\delta)\bar{\theta} + \delta\tilde{v}\right] + (1-h)\left\{a\left[(1-\delta)\underline{\theta} + \delta\tilde{v}\right] + (1-a)\delta v'\right\} \\
&= (1-\delta)\theta_E + \delta\tilde{v} + (1-a)(1-h)\left[\delta(v'-\tilde{v}) - (1-\delta)\underline{\theta}\right] \\
&\geq (1-\delta)\theta_E + \delta\tilde{v}, \text{ by agent IC.} \\
\implies \delta v - \delta\tilde{v} &\geq (1-\delta)\theta_E - (1-\delta)v \\
\implies v - \tilde{v} &\geq 2\epsilon, \text{ proving the claim.}
\end{aligned}$$

Claim 4: In any Pareto efficient equilibrium, there is (with probability 1) some period k such that the agent's continuation value is $\leq \omega$.

Proof: Combining Claims 1 and 3, there is a probability $\geq qh$ of the continuation value decreasing by at least ϵ in a given period. As $[0, \bar{v}]$ is bounded, the continuation value almost surely falls below ω at some time.

Claim 5: If the agent's continuation value is $v \leq \omega$ at some on-path history of a Pareto efficient equilibrium, his value is $\leq \omega$ forever.

Proof: Let b be the expected discounted number of bad projects from the current history. Appealing to Proposition 1, it must be that $b = 0$. Otherwise, we could replace the continuation play with a continuation equilibrium generating (v, ω) and preserve equilibrium. Therefore, by promise-keeping, it must be that every on-path future history (v^1, b^1) has $b^1 = 0$ as well. Again by Proposition 1, $v^1 \leq \omega$.

Together, Claims 4 and 5 deliver the theorem. □

As the relationship progresses, it systematically grows more conservative. From the agent's perspective, the early stages of the relationship are unequivocally better than the later stages; his degree of freedom declines over time. From the principal's perspective, the welfare implications are ambiguous; she wants to produce a high agent revenue as in the early stages, but she also wants a high yield as in the later stages.

5 Characterizing Equilibrium

We now know that the delegation relationship drifts toward conservatism in some (possibly distant) future, but this is only part of the story. How much productive delegation happens in the early, more flexible phase? When will bad projects be financed? While Theorem 1 tells us a lot about the trajectory of the relationship, it is silent on these more detailed questions about the nature of repeated delegation.

In this section, we explicitly describe the equilibrium value set of our repeated game—our main theorem. From this, we then derive some key implications for behavior in Pareto optimal equilibria.

5.1 The Equilibrium Value Set

First, we offer a complete characterization of the equilibrium value set.

The heart of the proof is the characterization of the equilibrium frontier B, formally carried out in the appendix. The overall structure of the argument proceeds as follows. First, allowing for public randomization guarantees us convexity of the equilibrium set frontier. As a consequence, whenever incentivizing picky project adoption by the agent, the principal optimally inflicts the minimum punishment possible. Next, because the principal has no private action or information, we show that the frontier is self-generating and private mixing unnecessary. Next, we show that initial freeze is inefficient for values above ω, and that bad project adoption is wasteful, except when used to provide very high agent values.

Theorem 2 (Equilibrium Value Set)**.** *There is a value* $\bar{v} > \omega$ *and a continuous function* $B : [0, \bar{v}] \to \mathbb{R}_+$ *such that* $\mathcal{E}^* = \{(v, b) \in [0, \bar{v}] \times \mathbb{R} : \ B(v) \leq b \leq \xi v\}$, *and:*

- $B|_{[0,\omega]} = 0$ *and* $B(\bar{v}) = (1-\delta)(1-h) + \delta B\left(\frac{v-(1-\delta)\omega}{\delta}\right)$.
- *On* $[\omega, \ \ \delta\bar{v} + (1-\delta)\omega]$, *B is strictly convex with*

$$B(v) = \delta\left[(1-h)B\left(\frac{v-(1-\delta)\omega}{\delta}\right) + hB\left(\frac{v-(1-\delta)\omega-(1-\delta)\underline{\theta}}{\delta}\right)\right].$$

- *On* $[\delta\bar{v} + (1-\delta)\omega, \ \ \bar{v}]$, *B is affine.*

It may at first appear surprising that one simple scalar equation can describe the full equilibrium set, even implicitly. This simple structure derives from the paucity of instruments at the principal's disposal. The key observation is that, with no ex-post monitoring of project types, the agent's binding incentive constraint completely pins down the law of motion for the agent's continuation value.

With $\mathcal{E}^*$ in hand, its Pareto frontier and the principal-optimal equilibrium value are immediate.

Corollary 2. *Let* $\bar{v}, B$ *be as in the statement of Theorem 2. There is a value* $v^* \in (\omega, \bar{v})$ *such that:*

- *The uniquely profit-maximizing equilibrium value profile is* $(v^*, B(v^*))$.
- *The Pareto frontier of* $\mathcal{E}^*$ *is* $\{(v, B(v)) : \ v^* \leq v \leq \bar{v}\}$.

Proof. Let $\pi^* : [0, \bar{v}] \to \mathbb{R}$ be given by $\pi^*(v) = \pi(v, B(v))$. Then $\operatorname{argmax}\pi^*$ is a subset of $[\omega, \bar{v}]$ because $\pi^*|_{[0,\omega]}$ is strictly increasing, nonempty because π^* is continuous, and a singleton (say v^*) because $\pi^*|_{[\omega,\bar{v}]}$ is strictly concave. The frontier characterization follows because π^* is increasing on $[0, v^*]$ and decreasing on $[v^*, \bar{v}]$.

It remains to verify that $v^* \neq \omega, \bar{v}$. It cannot be $\bar{v}$, since $\pi(v^*) \geq \pi^*(\omega) > 0 = \pi^*(\bar{v})$. Why can v^* not be ω? Well, by Lemma 12 in the Appendix, B is differentiable at ω with $B'(\omega) = 0$. Therefore, π^* is differentiable at ω with derivative > 0 and so cannot be maximized at ω. □

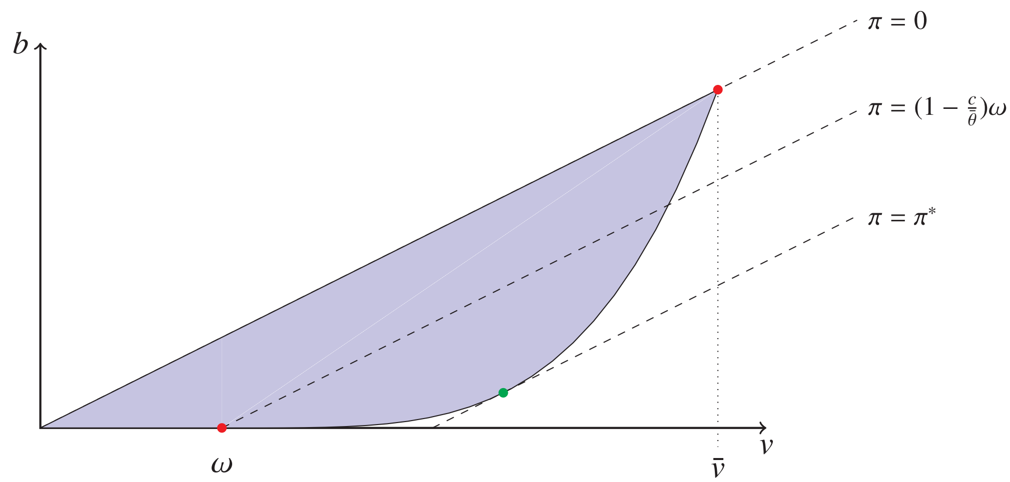

Figure 2: The solid line traces out the frontier B. The equilibrium value set is the convex region between B and the dashed zero-profit line. The dashed lines trace different isoprofits. The green dot highlights the uniquely principal-optimal vector.

5.2 Features of Equilibrium Behavior

We are now, with the equilibrium payoff set in hand, equipped to say more about the players' behavior under Pareto efficient equilibria.

Given an equilibrium, let $\{v_k,\ \hat{v}_k\}_{k=0}^{\infty}$ be the stochastic process of agent values, where v_k is before any public randomization that period and $\hat{v}_k$ is after. As Theorem 2 and its Corollary 2 show, every Pareto efficient equilibrium value (with agent value $v_0 \in [v^*, \bar{v}]$) can be implemented in an equilibrium with the following features (with probability 1 on the equilibrium path).

- **No superfluous randomization:** If $v_k \in [\omega, \delta\bar{v} + (1-\delta)\omega] \cup \{\bar{v}\}$, then $\hat{v}_k = v_k$. If $v_k \in (\delta\bar{v} + (1-\delta)\omega, \bar{v}]$, then $\hat{v}_k \in [\delta\bar{v} + (1-\delta)\omega, \bar{v}]$ too.

- **Good behavior with minimal punishment:** If $\hat{v}_k \in [\omega, \delta\bar{v} + (1-\delta)\omega)$, then the principal delegates, the agent adopts only good projects, and

$$v_{k+1} = \begin{cases} \frac{\hat{v}_k - (1-\delta)\omega}{\delta} & \text{if the agent doesn't adopt a project,} \\ \frac{\hat{v}_k - (1-\delta)\omega - (1-\delta)\underline{\theta}}{\delta} & \text{if the agent adopts a project.} \end{cases}$$

- **Gift-giving as a last resort:** If $\hat{v}_k > (1-\delta)\omega + \delta\bar{v}$, then the principal delegates, the agent adopts a project (irrespective of project type), and

$$v_{k+1} = \frac{\hat{v}_k - (1-\delta)\omega - (1-\delta)\underline{\theta}}{\delta}.$$

- **Absorbing punishment:** If $v_k \le \omega$, then $v_\ell, \hat{v}_\ell \le \omega$ for all $\ell \ge k$.

To gain some intuition as to why the above should be efficient, consider how the principal might like to provide different levels of revenue. The case of revenue $v \le \omega$ is simple: Proposition 1 tells us that the principal can provide said revenue efficiently, via aligned equilibrium. The case of $v = \bar{v}$ is straightforward too: to keep her current promise to the agent and offer no rewards tomorrow, the principal is forced to allow a project today, no questions asked. The described form of B between the two obtains from having the agent defend the principal's interests *whenever possible*. Reminiscent of **?**, costly incentive provision is backloaded as much as possible.[17]

Uniqueness As the next result shows, this backloading is a necessary feature of any Pareto optimal equilibrium. That is, the path of play implicit in Theorem 2—until the absorbing, conservative phase of the relationship is reached—is essentially[18] unique.

Corollary 3. *Every Pareto efficient equilibrium entails (i) no superfluous randomization, (ii) good behavior with minimal punishment, (iii) gift-giving as a last resort, and (iv) absorbing punishment.*

Proof. First recall that any Pareto efficient equilibrium will never have continuation value strictly above the graph of B at any on-path history. Indeed, replacing the

[17] With no private information, the costly incentive provision rewards the agent in **?** on-path. Here, the cost takes the form of (on-path) risk over the agent's continuation value.

[18] "Essentially" because immediate public randomization can be used over $(\delta\bar{v} + (1-\delta)\omega, \bar{v})$ but may not be required at values near $\bar{v}$, if δ is small enough. This region vanishes in the continuous time limit we consider in Section 6.

continuation value with one yielding the same agent value and strictly fewer bad projects will preserve equilibrium[19] and yield a Pareto improvement.

With this in mind, (iii) and (iv) follow directly from promise keeping, and (i) and half of (ii)—namely, minimum punishment—follow from strict convexity of $B|_{[\omega,\delta\bar{v}+(1-\delta)\omega]}$. We only need to show that the agent is picky whenever $\hat{v} \in (\omega, \delta\bar{v} + (1-\delta)\omega]$.

For values $\geq \omega$, freeze is never optimal Now, Theorem 2 tells us

$$\begin{aligned} D(v) \quad := \quad & \delta\left[(1-h)B\left(\frac{v-(1-\delta)\omega}{\delta}\right) + hB\left(\frac{v-(1-\delta)\omega-(1-\delta)\underline{\theta}}{\delta}\right)\right] \\ & -\left[(1-\delta)(1-h) + \delta B\left(\frac{v-(1-\delta)\omega-(1-\delta)\underline{\theta}}{\delta}\right)\right] \\ = \quad & (1-h)\delta\left[B\left(\frac{v-(1-\delta)\omega}{\delta}\right) - B\left(\frac{v-(1-\delta)\omega-(1-\delta)\underline{\theta}}{\delta}\right)\right] - (1-h)(1-\delta) \end{aligned}$$

is strictly increasing in $v \in [\omega, \delta\bar{v} + (1-\delta)\omega]$ (since B is convex, strictly so around $\frac{v-(1-\delta)\omega-(1-\delta)\underline{\theta}}{\delta}$) and globally ≤ 0 there, since asking the agent to be picky is optimal. Therefore, $D(v) < 0$ for $v < \delta\bar{v} + (1-\delta)\omega$. Therefore, unrestrained project adoption will not take place on the path of a Pareto efficient equilibrium while the agent's continuation value is in $[\omega, \delta\bar{v} + (1-\delta)\omega)$. Lastly, the agent must adopt all good projects there (by Theorem 1) and no bad projects (by Lemma 5). Point (ii) follows.

□

Bad Projects At the end Section 3, we saw that bad projects would, for some parameter values, be a necessary part of Pareto efficient equilibrium. Initially, one might have suspected that the choice of whether or not to employ bad projects to incentivize picky project adoption by the agent amounts to evaluating a profit tradeoff by the principal. Our main result, Theorem 2, tells us the principal faces no such tradeoff. Whenever a promise of future bad projects can be credible,[20] it is a necessary component of an optimal contract.

Corollary 4. *Every Pareto efficient equilibrium entails some bad projects.*

[19] For this logic, a special feature of our game is important: principal incentive compatibility amounts to a participation constraint.

[20] We have assumed they can, for expositional convenience, by imposing Assumption 3.

This follows directly from Corollary 2 (in particular, that $v^* > \omega$). Intuition for this result comes from considering equilibria as described in Corollary 3, with agent value $v = \omega + \epsilon$ strictly higher than, but very close to, the marginal value of search. The principal, with objective proportional to $\xi v - b$, benefits from the higher agent value $\omega + \epsilon > \omega$. Corollary 3 tells us that the cost this entails—a bounded discounted number of bad projects—will only accrue in an exponentially distant future. Thus the principal strictly prefers an equilibrium with $\omega + \epsilon$ to the aligned optimum.

The takeaway is straightforward: if bad projects can be credibly promised in equilibrium, then they unambiguously should be. The principal should offer them to the agent later to sustain better incentives today, prolonging the benefits of delegation.

6 Implementation: Dynamic Capital Budgeting

There are efficiency gains to be had from allowing bad projects, but one must carefully balance the principal's credibility constraint for it to remain an equilibrium. The implementation we discuss in this section, the **Dynamic Capital Budgeting** (DCB) contract, is an attempt to achieve this balance.

6.1 A Continuous Time Limit

For the purposes of this section, we find it convenient to work with a (heuristic) continuous time limit of our game in which the players interact very frequently, but good projects remain scarce. Letting the time between decisions, together with the proportion of good projects, vanish enables us to present our budget rule cleanly, in the language of calculus.

Suppose the players discount time $t \in \mathbb{R}_+$ at a rate of $r > 0$ and good projects arrive with Poisson rate $\eta > 0$. Assume bad projects are abundant, in the sense that (regardless of history) a bad project is available to adopt whenever a good project is not. The players meet at intervals of length $\Delta > 0$, so that period k corresponds to $t \in [k\Delta, (k+1)\Delta]$. Play within this interval is as follows:

- At time $k\Delta^+$, a public randomization device realizes, and the principal decides whether to delegate or freeze.

- Over the course of $(k\Delta, (k+1)\Delta)$, the agent sees the project arrival process. He then has in hand the best available project type:

$$\theta_k := \begin{cases} \bar{\theta} & \text{if at least one good project arrived during } (k\Delta, (k+1)\Delta), \\ \underline{\theta} & \text{otherwise.} \end{cases}$$

- If the principal delegated at time $k\Delta^+$, then the agent decides whether or not to adopt the project of type θ_k at time $(k+1)\Delta^-$. If the principal froze, the agent has no choice to make.

Rather than yielding flow payoffs, assume an initiated project of type θ provides the players a lump-sum revenue of θ, at a lump-sum cost[21] (to the principal) of c.

Up to a strategically irrelevant scaling of payoffs, the above continuous time game is equivalent to one in discrete time with parameters $\delta_\Delta := e^{-r\Delta}$, $h_\Delta := 1-e^{-\eta\Delta}$. Given fixed $\bar{\theta} > c > \underline{\theta} > 0$, we can invoke an appropriate lower bound[22] on $\frac{\eta}{r}$ to guarantee Assumptions 1, 2, and 3 hold for sufficiently small Δ.

Fix a sequence of period lengths $\Delta \to 0$ such that the equilibrium value set converges.[23] Let $\bar{v}_0$ be the highest agent value in the limit.

Adapting Theorem 2 tells us that $\bar{v}_0 > \omega_0$, where $\omega_0 := \eta(\bar{\theta} - \underline{\theta}) = \lim \frac{\delta_\Delta}{1-\delta_\Delta}\omega_\Delta$, the **marginal value of (instantaneous) search**. Corollary 3 then tells us that Pareto efficient equilibrium play has several interpretable features on-path in the limit. 'No superfluous randomization' tells us that, when the agent's continuation value is $\geq \omega_0$, no public randomization is currently used. 'Good behavior with minimal

[21]The analysis would not be changed if the benefit had a flow component but the cost were lump-sum, in which case θ would be interpreted as a present discounted value. If the cost were not lump-sum, on the other hand, the principal would face a new incentive constraint—to willingly continue to fund a costly project.

[22]Nothing is needed for Assumption 1. Then, $\frac{\eta}{r} > \frac{\underline{\theta}}{\bar{\theta}-\underline{\theta}}$ will imply Assumption 2. Given this, $\frac{\eta}{r}(\bar{\theta} - \underline{\theta})(1 - \frac{c}{\bar{\theta}}) > c - \underline{\theta}$ will yield Assumption 3.

[23]Some such sequence exists, where convergence is taken with respect to the Hausdorff metric. Indeed, each $\mathcal{E}^*_\Delta$ belongs to the compact set $\{\mathcal{E} \subseteq \bar{\mathcal{E}} : \mathcal{E} \text{ nonempty closed convex}\}$, where $\bar{\mathcal{E}} := \{(v, b) \in \mathbb{R}^2_+ : v \leq \eta\bar{\theta} + b\underline{\theta}, \ \pi(v, b) \geq 0\}$.

punishment' says that, if $v_t \in [\omega, \bar{v})$, then the principal delegates, the agent adopts only good projects, and

$$\begin{cases} v \text{ follows } \dot{v}_t = r(v_t - \omega_0) & \text{while the agent doesn't adopt a project,} \\ v \text{ jumps to } v_t - r\underline{\theta} & \text{if the agent adopts a project.} \end{cases}$$

'Gift-giving as a last resort' means that, when the continuation value hits $\bar{v}$, the principal continues to delegate to the agent, who immediately adopts a project (almost surely a bad one), and the agent's value jumps to $\bar{v} - r\underline{\theta}$. Lastly, 'absorbing punishment' tells us that, once the agent's continuation value is below ω_0, it stays there forever.

6.2 Defining the DCB contract

The DCB contract is characterized by a budget cap $\bar{x} \geq 0$ and initial budget balance $x \in [-1, \bar{x}]$, and consists of two regimes. At any time, players follow Controlled Budgeting or Capped Budgeting, depending on the agent's balance, x. The account balance can be understood as the number of projects the agent can initiate without immediately affecting the principal's delegation decisions.

Capped Budget $(x > 0)$
The account balance grows at the interest rate r as long as $x < \bar{x}$. Accrued interest is used to reward the agent for fiscal restraint. Since the opportunity cost of taking a project decreases in the account balance, the reward for diligence is increasing (exponentially) to maintain incentives. While there are funds in $\mathcal{A}$'s account, $\mathcal{P}$ fully delegates project choice to $\mathcal{A}$. However, every project that $\mathcal{A}$ initiates reduces the account balance to $x - 1$ (whether or not the latter is positive). Good projects being scarce, there are limits to how many projects the principal can credibly promise. When the balance is at the cap, the account can grow no further; accordingly, the agent takes a project immediately, yielding a balance of $\bar{x} - 1$.

Controlled Budget $(x \leq 0)$
The controlled budget regime is tailored to provide low revenue, low enough to be

feasibly provided in aligned equilibrium. When $x < 0$, the agent is over budget, and the principal punishes the agent—more severely the further over budget the agent is—with a freeze, restoring the balance to zero. The continuation contract when the balance is $x = 0$ is some optimal aligned equilibrium.

Definition. *The **Dynamic Capital Budgeting (DCB)** contract* $\sigma^{x,\bar{x}}$ *is as follows:*

1. *The **Capped Budget** regime:* $x > 0$.

   - *While* $x \in (0, \bar{x})$*:* $\mathcal{P}$ *delegates, and* $\mathcal{A}$ *takes any available good projects and no bad ones. If* $\mathcal{A}$ *initiates a project, the balance jumps from* x *to* $x - 1$*; if* $\mathcal{A}$ *doesn't take a project,* x *drifts according to* $\dot{x} = rx > 0$.
   - *When* x *hits* $\bar{x}$*:* $\mathcal{P}$ *delegates, and* $\mathcal{A}$ *takes a project immediately. If* $\mathcal{A}$ *picks up a project, the balance jumps from* $\bar{x}$ *to* $\bar{x} - 1$*; if* $\mathcal{A}$ *doesn't take a project, the balance remains at* $\bar{x}$.

2. *The **Controlled Budget** regime:* $x \leq 0$.

   - *If* $x \in [-1, 0)$*:* $\mathcal{P}$ *freezes for duration* $\frac{1}{r} \log \frac{\omega_0}{\omega_0 - \underline{\theta}|x|}$.
   - *After this initial freeze,* $\mathcal{P}$ *repeats the same policy forever: delegate until the next project, and then freeze for duration* $\bar{\tau} = \frac{1}{r} \log \frac{\omega_0}{\omega_0 - \underline{\theta}}$.

   In this regime, $\mathcal{A}$ *adopts every good project, and no bad projects.*

At the physics department's inception, the university allocates a budget of three hires, with a cap of ten. Over time, the physics department searches for candidates. Every time the department finds an appropriate candidate, it hires—and the provost rubber stamps it—spending from the agreed-upon budget. Figure 3 represents one possible realized path of the account balance over time.

The department finds two suitable candidates to hire in its first year; some interest having accrued, the department budget is now at two hires. After the first year, the department enters a dry spell: finding no suitable candidate for six years, the department hires no one. Due to accrued interest, the account budget has increased dramatically from two to over eight hires. In its ninth year, the account hits the cap and can grow no further. The department can immediately hire up to nine physicists and continue to search (with its remaining budget) or it can hire ten candidates and enact a regime change by the provost. The department chooses to hire one physics professor (irrespective of quality) immediately, and continue to search with a balance of nine.

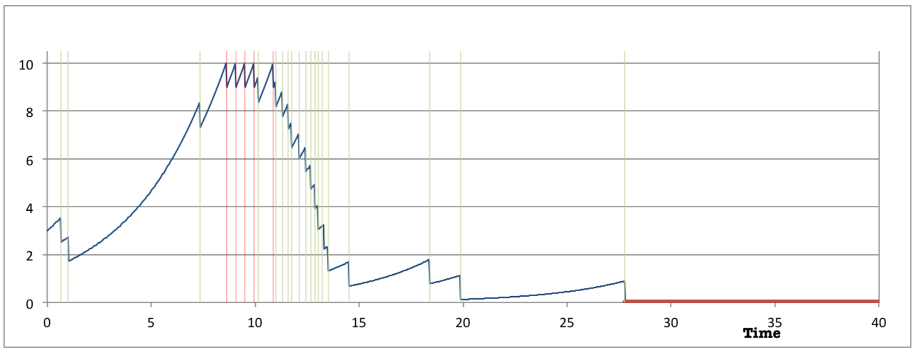

Figure 3: One realization of the balance's path under Controlled Budgeting (with $\bar{x} = 10$). Bad projects are clustered, and the account eventually runs dry.

Over the next few years, the department is flush and hires many professors. First, for three years, the department hits its cap several times, hiring many mediocre candidates. After its eleventh year, the department faces a lucky streak, finding many great physicists over the following years, bringing the budget to one hire. In the next twelve years, the department finds few candidates worth hiring. However, the interest accrual is so slow that the physics department still depletes its budget, in the twenty-eighth year. Throughout this initial phase, the department hires a total of twenty-four physics professors (much more than the account cap of ten).

At this point, the relationship changes permanently. After a temporary hiring freeze, the provost allows the department to resume its search, but follows any hire with a two-year hiring freeze. The relationship is now of a much more conservative character.

Notice that bad projects are clustered: the high balance $\bar{x} - 1$ just after a bad project means that the balance is likely to reach $\bar{x}$ again before the next arrival of a good project. So the next adopted project is likely bad. Given exponential growth, this effect is stronger the higher is the cap. In the Capped Budget regime, for a given account cap, the balance has non-monotonic profit implications. If the account runs low, there is an increased risk of imminently moving towards the low-revenue Controlled Budget regime. If the account runs high, the principal faces more bad projects in the near future. Observe that Controlled Budgeting is absorbing: once the balance falls low enough—which it eventually does—the agent will never take a bad project again.

Proposition 2. *Fixing an account cap and initial balance $\bar{x} > x > 0$, consider the Dynamic Capital Budget contract $\sigma^{x,\bar{x}}$.*

1. *Expected discounted revenue is $v(x) = \omega_0 + r\underline{\theta}x$.*
2. *Expected discounted number of bad projects is $b(x) = b^{\bar{x}}(x)$, uniquely determined by the delay differential equation*

$$(1 + \tfrac{\eta}{r})b(x) = \tfrac{\eta}{r}b(x-1) + xb'(x),$$

with boundary conditions:

$$\begin{aligned} b|_{(-\infty,0]} &= 0 \\ b(\bar{x}) - b(\bar{x}-1) &= 1. \end{aligned}$$

3. *$\sigma^{x,\bar{x}}$ is an equilibrium if and only if it exhibits nonnegative profit at the cap—that is,*

$$\bar{\pi}(\bar{x}) := \pi\Big(\omega_0 + r\underline{\theta}\bar{x},\ b(\bar{x})\Big) \geq 0.$$

Proof. The first point follows from substituting into the v promise-keeping constraint, and verifying that (by direct computation) $\sigma^{0,\bar{x}}$ yields revenue ω_0.

The second point follows from our work in Section 10.

For the third part, $v(x)-v(x-1) = [\omega_0+r\underline{\theta}x]-[\omega_0+r\underline{\theta}(x-1)] = r\underline{\theta}$ at every x, so that the agent is always indifferent between taking or leaving a bad project. Thus, $\sigma^{x,\bar{x}}$ is an equilibrium if and only if it satisfies principal participation after every history. Revenue is linear, and b is (by work in Section 10) convex. Therefore, profit is concave in x. So, profit is nonnegative for all on-path balances if and only if it is nonnegative at the top. □

6.3 Optimality of the DCB Contract

Finally we note that the DCB contract (with appropriate cap and initial balance) implements exactly the Pareto efficient equilibria. In this sense, our model gives an agency foundation to dynamic budgeting of capital expenditures. From the analysis of Subsection 6.1, the following is immediate.

Theorem 3. *Every Pareto efficient equilibrium payoff can be implemented as a DCB contract with cap* $\bar{x} := \frac{\bar{v}-\omega_0}{r\underline{\theta}}$ *and some initial balance.*

Adapting Corollary 3 gives a partial uniqueness result as well: every Pareto efficient equilibrium follows a two-regime structure, with Capped Budgeting (the cap computed as above) as its first regime.

6.4 Comparative Statics

In light of Theorem 3, a principal-optimal contract is characterized by two elements: how much freedom the principal can credibly give the agent (the cap), and how much freedom the principal chooses to initially give the agent (the initial balance). The first describes the equilibrium payoff set, while the second selects the principal-optimal contract therein. In this subsection, we ask how these features change as the environment the players face varies. As parameters of the model change, and the pool of projects becomes more valuable, the agent enjoys greater sovereignty, with both the balance cap and the optimal initial balance increasing.

Proposition 3. *For any profile of parameters that satisfy our Assumptions, define the account cap* $\bar{X}(\bar{\theta}, \underline{\theta}, c)$*, and the initial account balance* $X^*(\bar{\theta}, \underline{\theta}, c)$ *employed for principal-optimal equilibrium. Both functions are increasing in the revenue parameters* $\bar{\theta}, \underline{\theta}$*, and decreasing in the cost parameter* c*.*

The proof is in Section 11 of the appendix. We first analyze a slight increase in a revenue parameter and observe that the profit at the original cap increases. This implies that a slight cap increase maintains the principal's credibility, and it is now an equilibrium. This delivers the first half of our comparative statics result: the new equilibrium has a higher account cap, offering greater flexibility to the agent. As the account cap increases, the frontier of the equilibrium set gets flatter at each account balance. Moreover, as revenue parameters increase, the principal's isoprofit curves can only get steeper. Accordingly, the unique tangency between the equilibrium frontier and an isoprofit occurs at a higher balance. This delivers the remaining comparative statics result: the agent is also given more initial leeway. A similar analysis applies to a cost reduction.

Of particular interest is the role of $\underline{\theta}$ in determining the optimal DCB account structure. On the one hand, the principal suffers less from a bad project when $\underline{\theta}$ is higher; on the other, the agent is more tempted. We show that the former effect always dominates in determining how much freedom the principal optimally gives the agent.

7 Extensions

In this section, we briefly describe some extensions to our model. The proofs are straightforward and omitted.

Monetary Transfers

We maintain the assumption of limited liability: $\mathcal{A}$ cannot give $\mathcal{P}$ money. If $\mathcal{P}$ can reward $\mathcal{A}$'s fiscal restraint through direct transfers, one of two things happens: (i) nothing changes and money is not used; or (ii) money simply replaces bad projects as a reward if money is more cost-effective. Which is more efficient depends on the relative size of the marginal cost of allowing the agent to initiate bad projects[24] $(c-\underline{\theta})B'(\bar{v})$ and the marginal reward of doing so $\underline{\theta}$. If providing monetary incentives is optimal, a modified DCB contract is used. The cap is raised,[25] and the agent is paid a flow of cash whenever his balance is at the cap. This modified DCB contract is reminiscent of the optimal contract in **?**.

[24]This calculation is done using the B from our original model, as characterized in Theorem 2.

[25]The cap is raised to ensure zero profit with the new, more efficient incentivizing technology.

Permanent Termination

In many applications, being in a given relationship automatically entails delegating. If a client hires a lawyer, she delegates the choice of hours to be worked. To stop delegating is to terminate the relationship, giving both players zero continuation value.[26] That is, at any moment, the principal must choose between fully delegating and ending the game forever. At first glance, this constraint may seem an additional burden on the principal. However, given Theorem 1, we see that it changes nothing. Indeed, replacing any temporary freeze—guaranteed to happen only in the absorbing "$v \le \omega$" regime—with stochastic termination leaves payoffs and incentives unchanged.

Agent Replacement

Suppose the principal has a means to punish the agent without punishing herself: the principal can fire the agent and immediately hire a new one. The credibility of the threat of replacement takes us far from our leading examples: for instance, the state government cannot sever its relationship with one of its counties.

A fired agent gets a continuation payoff of zero.[27] Every time the principal hires a new agent, she proposes a new contract. Any inefficiency that the principal faces, as well as any interesting relationship dynamics in the contracts, vanish: in any equilibrium, the principal always delegates to the current agent, who, in turn, exercises fiscal restraint. Indeed, the principal can simply fully delegate to the current agent and replace him as soon as he takes a project. By Assumption 2, the agent finds it incentive compatible to take only good projects. This policy, featuring a far less dynamic relationship than in our model, yields first-best profits for the principal. Adapting Theorem 1, one can show that every equilibrium of this modified model entails a value $\le \omega$ for each agent.

[26]The agent could have a positive outside option. As long as it is below $\omega - (1-\delta)\underline{\theta}$, the same argument holds.

[27]Again, a positive agent outside option below $\omega - (1-\delta)\underline{\theta}$ would change nothing.

Commitment

If $\mathcal{P}$ has the ability to commit, she can offer $\mathcal{A}$ long-term rewards. In particular, she can offer him tenure (delegation forever) if he exerts fiscal restraint for a long enough time. With full commitment power, slight modifications of our argument show that $\bar{v}$ is the first-best revenue. **?** discuss this case more directly, using the methodology of **?**.

8 Final Remarks

In this paper, we have presented an infinitely repeated instance of the delegation problem. The agent will not represent the principal's interests without being offered dynamic incentives, while the principal cannot credibly commit to long-term rewards.

First, we speak to the evolution of the relationship under Pareto efficient equilibrium. Early on, the relationship is highly productive but low-yield: the agent adopts every good project, but some bad projects as well. The lack of principal commitment limits the magnitude of credible promises, making this phase a transient one. As the relationship matures, it is high-yield but less productive: the agent adopts only good projects, but some good opportunities go unrealized. In this sense, the relationship drifts toward conservatism.

Second, we solve for the exact form of an efficient intertemporal delegation rule, our Dynamic Capital Budget contract, which comprises two regimes. In the first regime, Capped Budgeting, the agent has an expense account, which grows at the interest rate so long as its balance is below its cap; the principal fully delegates, with every project being financed from the account. The agent takes every available good project; only when at the cap does he adopt projects indiscriminately. Eventually, the account runs dry, and the players transition to the second regime, Controlled Budgeting, wherein they play a (Pareto inefficient) aligned equilibrium. Not only is the DCB contract profit-maximizing, but it in fact traces out the whole equilibrium value set;[28] we note that the analysis and results apply at any fixed discount rate.[29]

While our main applications concern organizational economics outside of the firm, our results also speak to the canonical firm setting.[30] If the relationship between a firm and one of its departments proceeds largely via delegation, then we shed light on the dynamic nature of this relationship. In doing so, we provide a novel foundation for dynamic budgeting within the firm.

[28]More precisely, every equilibrium payoff is given by a mixture between the stage game equilibrium and a DCB contract.

[29]In particular, our analysis is not a folk theorem analysis.

[30]The conflict of interest in our model may reflect an empire-building motive on the part of a department, or it may be an expression of the **?** sales-maximization principle.

Appendix

This appendix provides formal proof for the results on which the main text draws. First, we provide a characterization of the equilibrium payoff set of the discrete-time game. Next, we prove several useful properties of the Delay Differential Equation which characterizes the frontier of the limit equilibrium set. Finally, we derive comparative statics results for the optimal cap and initial balance of a Dynamic Capital Budget contract.

9 APPENDIX: Characterizing Equilibrium Values

In the current section, we characterize the equilibrium value set in our discrete time repeated game. As in the main text, we find it convenient to study payoffs in terms of *agent value* and *bad projects*. Accordingly, for any strategy profile σ, we let

$$v(\sigma) = \mathbb{E}^{\sigma}\left[(1-\delta)\sum_{k=0}^{\infty}\delta^k \mathbf{1}_{\{\text{a project is adopted in period } k\}}\theta_k\right];$$
$$b(\sigma) = \mathbb{E}^{\sigma}\left[(1-\delta)\sum_{k=0}^{\infty}\delta^k \mathbf{1}_{\{\text{a project is adopted in period } k\}}1_{\theta_k=\underline{\theta}}\right].$$

Below, we will analyze the public perfect equilibrium (PPE) value set,

$$\mathcal{E}^* = \{(v(\sigma), b(\sigma)): \ \ \sigma \text{ is a PPE}\} \subseteq \mathbb{R}^2_+.$$

9.1 Self-Generation

To describe the equilibrium value set $\mathcal{E}^*$, we rely heavily on the machinery of **?**, called APS hereafter. To provide the players a given value $y = (v, b)$ from today onward, we factor it into a (possibly random) choice of what happens today, and what the continuation will be starting tomorrow. What happens today depends on the probability (p) that the principal delegates, the probability ($\bar{a}$) of project adoption if a project is good, and the probability ($\underline{a}$) of project adoption if a project is bad. The continuation values may vary based on what happens today: the principal may choose to freeze ($\check{y}$), the principal may delegate and agent may take a project ($\tilde{y}$), or the principal may delegate and agent may not take a project (y'). Since the principal doesn't observe project types, these are the only three public outcomes.

We formalize this factorization in the following definition and theorem.

Definition 2. *Given $Y \subseteq \mathbb{R}^2$:*

- *Say $y \in \mathbb{R}^2$ is **purely enforceable** w.r.t. Y if there exist $p, \bar{a}, a \in [0,1]$ and $\check{y}, \tilde{y}, y' \in Y$ such that:*[31]

  1. *(Promise keeping):*

$$\begin{aligned} y &= (1-p)\delta\check{y} + ph\left\{\bar{a}\left[(1-\delta)(\bar{\theta},0) + \delta\tilde{y}\right] + (1-\bar{a})\delta y'\right\} \\ &\quad + p(1-h)\left\{a\left[(1-\delta)(\underline{\theta},1) + \delta\tilde{y}\right] + (1-a)\delta y'\right\} \\ &= (1-p)\delta\check{y} + p\Big\{h\bar{a}\left[(1-\delta)(\bar{\theta},0) + \delta(\tilde{y}-y')\right] \\ &\quad + (1-h)a\left[(1-\delta)(\underline{\theta},1) + \delta(\tilde{y}-y')\right] + \delta y'\Big\}. \end{aligned}$$

  2. *(Incentive-compatibility):*

$$\begin{aligned} p \in \arg\max_{\hat{p}\in[0,1]} \quad & (1-\hat{p})\delta\check{\pi}(y) + \hat{p}\Big\{h\bar{a}\left[(1-\delta)(\bar{\theta}-c) + \delta[\pi(\tilde{y}) - \pi(y')]\right] \\ & + (1-h)a\left[(1-\delta)(\underline{\theta}-c) + \delta[\pi(\tilde{y}) - \pi(y')]\right] + \delta\pi(y') - \delta\pi(\check{y})\Big\}, \\ \bar{a} \in \arg\max_{\hat{a}\in[0,1]} \quad & \hat{a}\left\{(1-\delta)\bar{\theta} + \delta[v(\tilde{y}) - v(y')]\right\}, \\ a \in \arg\max_{\hat{a}\in[0,1]} \quad & \hat{a}\left\{(1-\delta)\underline{\theta} + \delta[v(\tilde{y}) - v(y')]\right\}. \end{aligned}$$

- *Say $y \in \mathbb{R}^2$ is **enforceable** w.r.t. Y if there exists a Borel probability measure μ on $\mathbb{R}^2$ such that*

  1. *$y = \int_{\mathbb{R}^2} \hat{y}\, d\mu(\hat{y})$.*
  2. *$\hat{y}$ is purely enforceable almost surely with respect to $\mu(\hat{y})$.*

- *Let $W(Y) := \{y \in \mathbb{R}^2 : \ y \text{ is enforceable with respect to } Y\}$.*

- *Say $Y \subseteq \mathbb{R}^2$ is **self-generating** if $Y \subseteq W(Y)$.*

Adapting methods from **?**, one can readily characterize $\mathcal{E}^*$ via self-generation, through the following collection of results.

[31] With a slight abuse of notation, for a given $y = (y_1, y_2) \in \mathbb{R}^2$, we will let $v(y) := y_1$.

Lemma 2. *Let W be as defined above.*

- *The set operator $W : 2^{\mathbb{R}^2} \longrightarrow 2^{\mathbb{R}^2}$ is monotone.*
- *Every bounded, self-generating $Y \subseteq \mathbb{R}^2$ is a subset of $\mathcal{E}^*$.*
- *$\mathcal{E}^*$ is the largest bounded self-generating set.*
- *$W(\mathcal{E}^*) = \mathcal{E}^*$.*
- *Let $Y_0 \subseteq \mathbb{R}^2$ be any bounded set with[32] $\mathcal{E}^* \subseteq W(Y_0) \subseteq Y_0$. Define the sequence $(Y_n)_{n=1}^{\infty}$ recursively by $Y_n := W(Y_{n-1})$ for each $n \in \mathbb{N}$. Then $\bigcap_{n=1}^{\infty} Y_n = \mathcal{E}^*$.*

9.2 A Cleaner Characterization

In light of the above, understanding the operator W will enable us to fully describe $\mathcal{E}^*$. That said, the definition of W is somewhat cumbersome. For the remainder of the current section, we work to better understand it.

Before doing anything else, we restrict attention to a useful domain for the map W.

Notation. *Let $\mathcal{Y} := \{Y \subseteq \mathbb{R}^2_+ : \ \vec{0} \in Y, \ \ Y$ is compact and convex, and $\pi|_Y \geq 0\}$.*

We need to work only with potential value sets in $\mathcal{Y}$. Indeed, the feasible set $\bar{\mathcal{E}}$ belongs to $\mathcal{Y}$, and it is straightforward to check that W takes elements of $\mathcal{Y}$ to $\mathcal{Y}$. Since $\mathcal{Y}$ is closed under intersections, we then know from the last bullet of the result above, Lemma 2, that $\mathcal{E}^* \in \mathcal{Y}$.

In seeking a better description of W, the following auxiliary definitions are useful.

Definition 3. *Given $a \in [0, 1]$:*

- *Say $y \in \mathbb{R}^2_+$ is a-**Pareto enforceable** w.r.t. Y if there exist $\tilde{y}, y' \in Y$ such that:*

  1. *(Promise keeping):*

$$y = h\left[(1-\delta)(\bar{\theta}, 0) + \delta(\tilde{y} - y')\right] + (1-h)a\left[(1-\delta)(\underline{\theta}, 1) + \delta(\tilde{y} - y')\right] + \delta y'.$$

[32] This can be ensured, for instance, by letting Y_0 contain the feasible set, scaled by $\frac{1}{1-\delta}$.

2. *(Agent incentive-compatibility):*

$$1 \in \arg\max_{\hat{a}\in[0,1]} \quad \hat{a}\left\{(1-\delta)\bar{\theta} + \delta[v(\tilde{y}) - v(y')]\right\},$$
$$a \in \arg\max_{\hat{a}\in[0,1]} \quad \hat{a}\left\{(1-\delta)\underline{\theta} + \delta[v(\tilde{y}) - v(y')]\right\}.$$

3. *(Principal participation):* $\pi(y) \geq 0$.

- *Let* $W_a(Y) := \{y \in \mathbb{R}^2_+ : \ y \text{ is } a\text{-Pareto enforceable w.r.t. } Y\}$.
- *Let* $W_f(Y) := \delta Y$.
- *Let* $\hat{W}(Y) := W_f(Y) \cup \bigcup_{a\in[0,1]} W_a(Y)$*. If* Y *is compact, then so is* $\hat{W}(Y)$.[33]

The set $\hat{W}(Y)$ captures the enforceable (without public randomizations) values w.r.t. Y if:

1. The principal uses a pure strategy.
2. We relax principal IC to a participation constraint.
3. If the principal delegates and the project is good, then the agent takes the project.

The following proposition shows that, for the relevant $Y \in \mathcal{Y}$, it is without loss to focus on $co\hat{W}$ instead of W. The result is intuitive. The first two points are without loss because the principal's choices are observable. Toward (1), her private mixing can be replaced with public mixing with no effect on $\mathcal{A}$'s incentives. Toward (2), if the principal faces nonnegative profits with any pure action, she can be incentivized to take said action with stage Nash (min-max payoffs) continuation following the other choice. Toward (3), the agent's private mixing isn't (given (2)) important for the principal's IC, and so we can replace it with public mixing between efficient (i.e. no good project being passed up) first-stage play and an initial freeze.

Lemma 3. *Under Assumption 1, if* $Y \in \mathcal{Y}$*, then* $W(Y) = co\hat{W}(Y)$.

Proof. First, notice that $\delta Y \subseteq W(Y) \cap co\hat{W}(Y)$. It is a subset of the latter by construction, and of the former by choosing $\tilde{y} = y' = \vec{0}$, $p = 0, \bar{a} = a = 1$, and letting $\check{y}$ range over Y.

[33] Indeed, it is the union of δY and a projection of the compact set $\{(a, y) \in [0, 1] \times \mathbb{R}^2 : y \text{ is } a\text{-Pareto enforceable w.r.t. } Y\}$.

Take any $y \in \hat{W}(Y)$ that isn't in δY. So y is a-Pareto enforceable w.r.t. Y for some $a \in [0,1]$, say witnessed by $\tilde{y}, y' \in Y$. Letting $p = 1$, $\bar{a} = 1$, and $\check{y} = \vec{0} \in Y$, it is immediate that $p, \bar{a}, a \in [0,1]$ and $\check{y}, \tilde{y}, y' \in Y$ witness y being purely enforceable w.r.t. Y. Therefore, $y \in W(Y)$. So $\hat{W}(Y) \subseteq W(Y)$. The latter being convex, $co\hat{W}(Y) \subseteq W(Y)$ as well.

Take any extreme point y of $W(Y)$ which isn't in δY. Then y must be purely enforceable w.r.t. Y, say witnessed by $p, \bar{a}, a \in [0,1]$ and $\check{y}, \tilde{y}, y' \in Y$. First, if $p\bar{a} = 0$, then[34]

$$y = (1-p)\delta\check{y} \in co\{\vec{0}, \delta\check{y}\} \subseteq \delta Y \subseteq \hat{W}(Y).$$

Now suppose $p\bar{a} > 0$, and define[35] $a^p := \dfrac{a}{\bar{a}}$ and

$$y^p := h\left[(1-\delta)(\bar{\theta}, 0) + \delta(\tilde{y} - y')\right] + (1-h)a^p\left[(1-\delta)(\underline{\theta}, 1) + \delta(\tilde{y} - y')\right] + \delta y'.$$

Observe that $\tilde{y}, y'$ witness $y^p \in W_{a^p}(Y)$:

1. Promise keeping follows immediately from the definition of y^p.

2. Agent IC follows from agent IC in enforcement of y, and from the fact that incentive constraints are linear in action choices. As $\bar{a} > 0$ was optimal, $\bar{a}^p = 1$ is optimal here as well.

3. Principal participation follows from principal IC in enforcement of y, and from the fact that $\pi(\check{y}) \geq 0$ because $\pi|_Y \geq 0$.

Therefore $y^p \in W_{a^p}(Y)$, from which it follows that

$$y = (1-p)\delta\check{y} + p\bar{a}y^p \in co\{\delta\check{y}, y^p, \vec{0}\} \subseteq \hat{W}(Y).$$

As every extreme point of $W(Y)$ belongs to $\hat{W}(Y)$, all of $W(Y)$ belongs to the closed convex hull of $\hat{W}(Y)$, which is just $co\hat{W}(Y)$.[36] □

In view of the above proposition, we now only have to consider the much simpler map $\hat{W}$. As the following lemma shows, we can even further simplify, by showing that there is never a need to offer excessive punishment. That is, it is without loss to (1) make the agent's

[34] If $p\bar{a} = 0$, then either $p = 0$ or $\bar{a} = 0$. If $\bar{a} = 0$, then agent IC implies $a = 0$. So either $p = 0$ or $a = \bar{a} = 0$; in either case, promise keeping then implies $y = (1-p)\delta\check{y}$.

[35] Since $a \leq \bar{a}$ by IC, we know $a^p \in [0,1]$.

[36] The disappearance of the qualifier "closed" comes from Carathéodory's theorem, since $\hat{W}(Y)$ is compact in Euclidean space.

IC constraint (to resist bad projects) bind if he is being discerning, and (2) not respond to the agent's choice if he is being indiscriminate.

Lemma 4. *Suppose Assumption 1 holds. Fix $a \in [0,1]$, $Y \in \mathcal{Y}$, and $y \in \mathbb{R}^2$:*
Suppose $a < 1$. Then $y \in W_a(Y)$ if and only if there exist $\tilde{z}, z' \in Y$ such that:

1. *(Promise keeping):*

$$y \quad = \quad h\left[(1-\delta)(\bar{\theta},0)+\delta(\tilde{z}-z')\right]+(1-h)a\left[(1-\delta)(\underline{\theta},1)+\delta(\tilde{z}-z')\right]+\delta z'.$$

2. *(Agent **exact** incentive-compatibility):*

$$\delta[v(z')-v(\tilde{z})] = (1-\delta)\underline{\theta}.$$

3. *(Principal participation): $\pi(y) \geq 0$.*

Suppose $a = 1$. Then $y \in W_a(Y)$ if and only if there exists $z' \in Y$ such that:

1. *(Promise keeping):*

$$y = (1-\delta)[h(\bar{\theta},0)+(1-h)(\underline{\theta},1)]+\delta z'$$

2. *(Principal participation): $\pi(y) \geq 0$.*

Proof. In the first case, the "if" direction is immediate from the definition of W_a. In the second, it is immediate once we apply the definition of W_1 with $\tilde{z} = z'$. Now we proceed to the "only if" direction.

Consider any $y \in W_a(Y)$, with $\tilde{y}, y'$ witnessing a-Pareto enforceability. Define

$$\bar{y} := [h + a(1-h)]\tilde{y} + (1-h)(1-a)y' \in Y.$$

So $\bar{y}$ is the on-path expected continuation value.

In the case of $a < 1$, define

$$\begin{aligned} q &:= \frac{(1-\delta)\underline{\theta}}{\delta[v(y')-v(\tilde{y})]} \quad (\in [0,1], \text{ by IC}) \\ \tilde{z} &:= (1-q)\bar{y}+q\tilde{y} \\ z' &:= (1-q)\bar{y}+qy'. \end{aligned}$$

By construction, $\delta[v(z') - v(\tilde{z})] = (1-\delta)\underline{\theta}$, as desired.[37]

In the case of $a = 1$, let $z' := \bar{y}$ and $\tilde{z} := z'$.

Notice that $\tilde{z}, z'$ witness $y \in W_a(Y)$. Promise keeping comes from the definition of $\bar{y}$, principal participation comes from the hypothesis that $W_a(Y) \ni y$, and IC (exact in the case of $a < 1$) comes by construction. □

In the first part of the lemma, $\delta[v(y') - v(\tilde{y})] \in (1-\delta)[\underline{\theta}, \bar{\theta}]$ has been replaced with $\delta[v(z') - v(\tilde{z})] = (1-\delta)\underline{\theta}$. That is, it is without loss to make the agent's relevant incentive constraint—to avoid taking bad projects—bind. This follows from the fact that $Y \supseteq co\{\tilde{y}, y'\}$. The second part of the lemma says that, if the agent isn't being at all discerning, nothing is gained from disciplining him.

The above lemma has a clear interpretation: without loss of generality, the principal uses the minimal possible punishment. The lemma also yields the following:

Lemma 5. *Suppose Assumption 1 holds. Suppose $a \in (0,1)$, $Y \in \mathcal{Y}$, and $y \in W_a(Y)$. Then there is some $y^* \in W_0$ such that*

$$v(y^*) = v(y) \text{ and } b(y^*) < b(y).$$

That is, $y_1^* = y_1$ and $y_2^* < y_2$.

Proof. Appealing to Lemma 4, there exist $\tilde{z}, z' \in Y$ such that:

1. (Promise keeping):

$$y = h\left[(1-\delta)(\bar{\theta}, 0) + \delta(\tilde{z} - z')\right] + (1-h)a\left[(1-\delta)(\underline{\theta}, 1) + \delta(\tilde{z} - z')\right] + \delta z'$$

2. (Agent **exact** incentive-compatibility):

$$\delta[v(z') - v(\tilde{z})] = (1-\delta)\underline{\theta}$$

3. (Principal participation): $\pi(y) \geq 0$.

Given agent exact IC, we know $v(z') > v(\tilde{z})$. Let $\tilde{z}^* := \left(v(\tilde{z}), \min\left\{b(\tilde{z}), \frac{v(\tilde{z})}{v(z')} b(z')\right\}\right)$. As either $\tilde{z}^* = \tilde{z}$ or $\tilde{z}^* \in co\{\vec{0}, z'\}$, we have $z^* \in Y$.

[37] In the case of $a \in (0,1)$, $q = 1$ (by agent IC), so that $\tilde{z} = \tilde{y}$ and $z' = y'$. The real work was needed for the case of $a = 0$.

Let $y^* := h\left[(1-\delta)(\bar{\theta}, 0) + \delta(\tilde{z}^* - z')\right] + \delta z'$. Then

$$\begin{aligned} v(y) - v(y^*) &= (1-h)a\left\{(1-\delta)\underline{\theta} + \delta[v(\tilde{z}^*) - v(z')]\right\} - h\delta[v(\tilde{z}^*) - v(\tilde{z})] \\ &= (1-h)a\left\{(1-\delta)\underline{\theta} + \delta[v(\tilde{z}) - v(z')]\right\} - h\delta 0 \\ &= 0, \end{aligned}$$

while

$$\begin{aligned} b(y) - b(y^*) &= (1-h)a\{(1-\delta) + \delta[b(\tilde{z}) - b(z')]\} - h\delta[b(\tilde{z}^*) - b(\tilde{z})] \\ &= (1-h)a\{(1-\delta) + \delta[b(\tilde{z}^*) - b(z')] + \delta[b(\tilde{z}) - b(\tilde{z}^*)]\} - h\delta[b(\tilde{z}^*) - b(\tilde{z})] \\ &= (1-h)a\{(1-\delta) + \delta[b(\tilde{z}^*) - b(z')]\} + [h + (1-h)a]\delta[b(\tilde{z}) - b(\tilde{z}^*)] \\ &\geq (1-h)a \\ &> 0. \end{aligned}$$

Now, notice that $\tilde{z}^*, z'$ witness $y^* \in W_0(Y)$. Promise keeping holds by fiat, agent IC holds because $v(\tilde{z}^*) = v(\tilde{z})$ by construction, and principal participation follows from

$$\pi(y^*) - \pi(y) = -\pi(0, b(y) - b(y^*)) > 0.$$

□

The above lemma is a strong bang-bang result. It isn't just sufficient to restrict attention to equilibria with no private mixing; it is necessary too. Any equilibrium in which the agent mixes on-path is Pareto dominated.

9.3 Self-Generation for Frontiers

Through Lemmata 3, 4, and 5, we greatly simplified analysis of the APS operator W applied to the relevant value sets. In the current section, we profit from that simplification in characterizing the efficient frontier of $\mathcal{E}^*$. Before we can do that, however, we have to make a small investment in some new definitions. We then translate the key results of the previous subsection into results about the efficient frontier of the equilibrium set.

Notation. *Let $\mathcal{B}$ denote the space of functions $B : \mathbb{R} \longrightarrow \mathbb{R}_+ \cup \{\infty\}$ such that: (1) B is convex, (2) $B(0) = 0$, (3) B's proper domain $dom(B) := B^{-1}(\mathbb{R})$ is a compact subset of $\mathbb{R}_+$, (4) B is continuous on $dom(B)$, and (5) $\pi(v, B(v)) \geq 0$ for every $v \in dom(B)$.*

Just as $\mathcal{Y}$ is the relevant space of value sets, $\mathcal{B}$ is the relevant space of frontiers of value sets.

Notation. *For each $Y \in \mathcal{Y}$, define the **efficient frontier function** of Y:*

$$\begin{aligned} B_Y : \mathbb{R} &\longrightarrow \mathbb{R}_+ \cup \{\infty\} \\ v &\longmapsto \min\{b \in \mathbb{R}_+ : \ (v, b) \in Y\} \end{aligned}$$

It is immediate that for $Y \in \mathcal{Y}$, the function B_Y belongs to $\mathcal{B}$.

Notation. *Define the following functions:*

$$\begin{aligned} T : \mathcal{B} &\longrightarrow \mathcal{B} \\ \hat{B} &\longmapsto B_{W\left(co[graph(\hat{B})]\right)} = B_{co\hat{W}\left(co[graph(\hat{B})]\right)}, \\ T_f : \mathcal{B} &\longrightarrow \mathcal{B} \\ \hat{B} &\longmapsto B_{W_f\left(co[graph(\hat{B})]\right)} = B_{\delta\left(co[graph(\hat{B})]\right)}, \\ \textit{and for } 0 \le a \le 1, \quad T_a : \mathcal{B} &\longrightarrow \mathcal{B} \\ \hat{B} &\longmapsto B_{W_a\left(co[graph(\hat{B})]\right)}. \end{aligned}$$

These objects are not new. The map T [resp. T_f, T_a] is just a repackaging of W [resp. W_f, W_a], made to operate on frontiers of value sets, rather than on value sets themselves.

As it turns out, we really only need Lemmata 3, 4, and 5 to simplify our analysis of T, which in turn helps us characterize the efficient frontier of $\mathcal{E}^*$. We now proceed along these lines.

The following lemma is immediate from the definition of the map $Y \mapsto B_Y$.

Lemma 6. *If $\{Y_i\}_{i \in \mathbb{I}} \subseteq \mathcal{Y}$, then $B_{\overline{co}[\bigcup_{i \in \mathbb{I}} Y_i]}$ is the convex lower envelope*[38] *of* $\inf_{i \in \mathbb{I}} B_{Y_i}$.

The following proposition is the heart of our main characterization of the set $\mathcal{E}^*$'s frontier. It amounts to a complete description of the behavior of T.

Lemma 7. *Suppose Assumption 1 holds. Fix any $B \in \mathcal{B}$ and $v \in \mathbb{R}$. Then:*

1. $TB = cvx\left[\min\left\{T_f B, T_0 B, T_1 B\right\}\right]$.

[38] The **convex lower envelope** of a function $\check{B}$ is $\text{cvx}\check{B}$,the largest convex upper-semicontinuous function below it. Equivalently, $\text{cvx}\check{B}$ is the pointwise supremum of all affine functions below $\check{B}$.

2. *For* $i \in \{f, 0, 1\}$,

$$T_i B(v) = \begin{cases} \check{T}_i B(v) & \text{if } \pi\left(v, \check{T}_i^{\Delta} B(v)\right) \geq 0 \\ \infty & \text{otherwise,} \end{cases}$$

where

$$\begin{aligned}
\check{T}_f B(v) &:= \delta B\left(\frac{v}{\delta}\right) \\
\check{T}_0 B(v) &:= \delta\left[hB\left(\frac{v-(1-\delta)\theta_E}{\delta}\right) + (1-h)B\left(\frac{v-(1-\delta)[\theta_E - \underline{\theta}]}{\delta}\right)\right] \\
&= \delta\left[hB\left(\frac{v-(1-\delta)\omega-(1-\delta)\underline{\theta}}{\delta}\right) + (1-h)B\left(\frac{v-(1-\delta)\omega}{\delta}\right)\right] \\
\check{T}_1 B(v) &:= (1-h)(1-\delta) + \delta B\left(\frac{v-(1-\delta)\theta_E}{\delta}\right) \\
&= (1-h)(1-\delta) + \delta B\left(\frac{v-(1-\delta)\omega-(1-\delta)\underline{\theta}}{\delta}\right)
\end{aligned}$$

Proof. That $TB = \text{cvx}\left[\min\left\{T_f B, \ \inf_{a\in[0,1]} T_a B\right\}\right]$ is a direct application of Lemma 6. Then, appealing to Lemma 5, $T_a B \geq T_0 B$ for every $a \in (0, 1)$. This proves the first point.

In what follows, let $Y := co[\text{graph}(B)]$ so that $TB = B_{W(Y)}$.

- Consider any $y \in W_0(Y)$:
  Lemma 4 delivers $\tilde{z}, z' \in Y$ such that

$$\begin{aligned}
y &= h\left[(1-\delta)(\bar{\theta}, 0) + \delta(\tilde{z} - z')\right] + \delta z', \\
(1-\delta)\underline{\theta} &= \delta[v(z') - v(\tilde{z})].
\end{aligned}$$

  Rewriting with $\tilde{z} = (\tilde{v}, \tilde{b})$ and $z' = (v', b')$, and rearranging yields:

$$\begin{aligned}
(1-\delta)\underline{\theta} &= \delta[v' - \tilde{v}] \\
(v, b) &= h\left[(1-\delta)(\bar{\theta}, 0) + \delta(\tilde{v} - v', \ \tilde{b} - b')\right] + \delta(v', b') \\
&= h\left((1-\delta)(\bar{\theta} - \underline{\theta}), \ \delta[\tilde{b} - b']\right) + \delta(v', b') \\
&= \left(\theta_E - \underline{\theta} + \delta v', \ \ h\delta\tilde{b} + (1-h)\delta b'\right)
\end{aligned}$$

Solving for the agent values yields

$$v' = \frac{v - (1-\delta)[\theta_E - \underline{\theta}]}{\delta} \text{ and } \tilde{v} = v' - \delta^{-1}(1-\delta)\underline{\theta} = \frac{v - (1-\delta)\theta_E}{\delta}.$$

So given any $v \in \mathbb{R}_+$:

$$\begin{aligned} T_0 B(v) &= \inf_{b,\tilde{b},b'} b \\ &\text{s.t.} \quad \pi(v,b) \geq 0, \quad b = \delta\left[h\tilde{b} + (1-h)b'\right], \\ &\quad \text{and} \left(\frac{v - (1-\delta)[\theta_E - \underline{\theta}]}{\delta}, \; b'\right), \left(\frac{v - (1-\delta)\theta_E}{\delta}, \; \tilde{b}\right) \in Y \\ &= \inf_{b,\tilde{b},b'} b = \delta\left[h\tilde{b} + (1-h)b'\right] \\ &\text{s.t.} \quad \pi(v,b) \geq 0 \text{ and } \left(\frac{v - (1-\delta)[\theta_E - \underline{\theta}]}{\delta}, \; b'\right), \left(\frac{v - (1-\delta)\theta_E}{\delta}, \; \tilde{b}\right) \in Y \\ &= \begin{cases} b = \delta\left[hB\left(\frac{v-(1-\delta)\theta_E}{\delta}\right) + (1-h)B\left(\frac{v-(1-\delta)[\theta_E - \underline{\theta}]}{\delta}\right)\right] & \text{if } \pi(v,b) \geq 0, \\ \infty & \text{otherwise.} \end{cases} \end{aligned}$$

- Consider any $y \in W_1(Y)$:
  Lemma 4 now delivers $z' = (v', b') \in Y$ such that

$$y = (1-\delta)[h(\bar{\theta}, 0) + (1-h)(\underline{\theta}, 1)] + \delta z',$$

which can be rearranged to

$$(v, b) = \left((1-\delta)\theta_E + \delta v', \;\; (1-h)(1-\delta) + \delta b'\right).$$

So given any $v \in \mathbb{R}_+$:

$$\begin{aligned}
T_1 B(v) &= \inf_{b,b'} b \\
&\text{s.t.} \quad \pi(v,b) \geq 0, \quad b = (1-h)(1-\delta) + \delta b', \text{ and } \left(\frac{v - (1-\delta)\theta_E}{\delta}, \; b' \right) \in Y \\
&= \inf_{\tilde{b},b'} b = (1-h)(1-\delta) + \delta b' \\
&\text{s.t.} \quad \pi(v,b) \geq 0 \text{ and } \left(\frac{v - (1-\delta)\theta_E}{\delta}, \; b' \right) \in Y \\
&= \begin{cases} b = (1-h)(1-\delta)\delta B\left(\frac{v-(1-\delta)\theta_E}{\delta}\right) & \text{if } \pi(v,b) \geq 0, \\ \infty & \text{otherwise.} \end{cases}
\end{aligned}$$

- Lastly, given any $v \in \mathbb{R}_+$:

$$\begin{aligned}
T_f B(v) &= \inf_{b,b'} b \\
&\text{s.t.} \quad \pi(v,b) \geq 0, \quad b = \delta b', \text{ and } \left(\frac{v}{\delta}, b' \right) \in Y \\
&= \begin{cases} b = \delta B\left(\frac{v}{\delta}\right) & \text{if } \pi(v,b) \geq 0, \\ \infty & \text{otherwise.} \end{cases}
\end{aligned}$$

□

9.4 The Efficient Frontier

In this subsection, we characterize the frontier $B_{\mathcal{E}^*}$ of the equilibrium value set. We first translate APS's self-generation to the setting of frontiers. This, along with Lemma 7 delivers our Bellman equation, Corollary 5. Then, in Theorem 1, we characterize the discrete time equivalent of the Aligned Optimal Budget. Finally, in Lemma 11, we fully characterize the frontier $B_{\mathcal{E}^*}$.

Lemma 8. *Suppose Assumption 1 holds. Suppose $Y \in \mathcal{Y}$ with $W(Y) = Y$. Then $TB_Y = B_Y$.*

Proof. First, because W is monotone, $W(co[\text{graph}(B_Y)]) \subseteq W(Y) = Y$. Thus the efficient frontier of the former is higher than that of the latter. That is, $TB_Y \geq B_Y$.

Now take any $v \in \text{dom}(B_Y)$ such that $y := (v, B_Y(v))$ is an extreme point of Y. We want to show that $TB_Y(v) \leq B_Y(v)$.

By Lemma 3, $y \in W_f(Y) \cup \bigcup_{a\in[0,1]} W_a(Y)$.

- If $y \in W_f(Y)$, then $\frac{y}{\delta}, \vec{0} \in Y$, so that the extreme point y must be equal to $\vec{0}$. But in this case, $TB_Y(v) = TB_Y(0) = 0 = B_Y(0) = B_Y(v)$.

- If $y \in W_a(Y)$ for some $a \in [0,1]$, say witnessed by $\tilde{y}, y' \in Y$, then let

$$\begin{aligned}
\tilde{z} &:= (v(\tilde{y}), B_Y(v(\tilde{y}))) \\
\tilde{z}' &:= (v(y'), B_Y(v(y'))) \\
z &:= h\left[(1-\delta)(\bar{\theta}, 0) + \delta(\tilde{z} - z')\right] + (1-h)a\left[(1-\delta)(\underline{\theta}, 1) + \delta(\tilde{z} - z')\right] + \delta z'.
\end{aligned}$$

  Then

$$\begin{aligned}
b(z) &= (1-h)(1-\delta)a + [h + (1-h)a]\delta B_Y(v(\tilde{y})) + (1-h)(1-a)\delta B_Y(v(y')) \\
&\leq (1-h)(1-\delta)a + [h + (1-h)a]\delta b(\tilde{y}) + (1-h)(1-a)\delta b(y') \\
&= b(y) = B_Y(v),
\end{aligned}$$

  and $\tilde{z}, z'$ witness $z \in W_a(co[\text{graph}(B_Y)])$. In particular, $TB_Y(v) = TB_Y(v(z)) \leq b(z) \leq B_Y(v)$

Next, consider any $v \in \text{dom}(B_Y)$. There is some probability measure μ on the extreme points of Y such that $(v, B_Y(v)) = \int_Y y \, d\mu(y)$. By minimality of $B_Y(v)$, it must be that $y \in \text{graph}(B_Y)$ a.s.-$\mu(y)$. So letting μ_1 be the marginal of μ on the first coordinate, $(v, B_Y(v)) = \int_{v(Y)} (u, B_Y(u)) \, d\mu_1(u)$, so that

$$B_Y(v) = \int_{v(Y)} B_Y \, d\mu_1 \geq \int_{v(Y)} TB_Y \, d\mu_1 \geq TB_Y(v),$$

where the last inequality follows from Jensen's theorem.

This completes the proof.

□

The Bellman equation follows immediately.

Corollary 5. *Under Assumption 1, $B := B_{\mathcal{E}}$ solves the Bellman equation $B = cvx\left[\min\left\{T_f B,\ T_0 B,\ T_1 B\right\}\right]$.*

Aligned Optimal Equilibrium In line with the main text, we now proceed to characterize the payoffs attainable in equilibria with no bad projects.

Lemma 9. *Suppose Assumption 1 holds.*

1. *There exist productive aligned equilibria if and only if Assumption 2 holds.*
2. *Every aligned equilibrium generates revenue* $\leq \omega$.
3. *Given Assumption 2,* $(\omega, 0) \in \mathcal{E}^*$.

Proof. We proceed in reverse, invoking Lemma 7 throughout.

For (3), define $B \in \mathcal{B}$ via $B(v) = 0$ for $v \in [0, \omega]$ and $B(v) = \infty$ otherwise. Notice that, $TB(\omega) \leq T_0 B(\omega) = \delta h B(\frac{\omega-(1-\delta)\theta_E}{\delta}) + \delta(1-h)B(\frac{\omega-(1-\delta)\omega}{\delta}) = 0$, where the second equality holds by Assumption 2. Because TB is convex, B is self-generating, and thus $B \geq B_{\mathcal{E}^*}$. (3) follows.

We now proceed to verify (2), i.e. that $\hat{v} > \omega$ implies that $(0, \hat{v}) \notin \mathcal{E}^*$.
Suppose $v > \omega$ has $B_{\mathcal{E}^*}(v) = 0$. Then $B_{\mathcal{E}^*}|_{[0,v]} = 0$, and

$$0 = B_{\mathcal{E}^*}(v) = TB_{\mathcal{E}^*}(v) = \min\{T_f B_{\mathcal{E}^*}(v), T_0 B_{\mathcal{E}^*}(v), T_1 B_{\mathcal{E}^*}(v)\}.$$

Notice that $TB_{\mathcal{E}^*}(v) \neq T_1 B_{\mathcal{E}^*}(v)$ as the latter is > 0. If $TB_{\mathcal{E}^*}(v) = T_f B_{\mathcal{E}^*}(v)$, then since $B_{\mathcal{E}^*}$ is increasing

$$B_{\mathcal{E}^*}\left(v + \frac{1-\delta}{\delta}(v-\omega)\right) \leq B_{\mathcal{E}^*}\left(v + \frac{1-\delta}{\delta}v\right) = \delta^{-1} T_f B_{\mathcal{E}^*}(v) = 0.$$

Finally, if $TB_{\mathcal{E}^*}(v) = T_0 B_{\mathcal{E}^*}(v)$, then[39]

$$0 = B_{\mathcal{E}^*}\left(\frac{v-(1-\delta)\omega}{\delta}\right) = \delta B_{\mathcal{E}^*}\left(v + \frac{1-\delta}{\delta}(v-\omega)\right).$$

So either way, $B_{\mathcal{E}^*}\left(v + \frac{1-\delta}{\delta}(v-\omega)\right) = 0$ too.
Now, if $\hat{v} > \omega$ with $(0, \hat{v}) \in \mathcal{E}^*$, then applying the above inductively yields a sequence[40] $v_n \to \infty$ on which $B_{\mathcal{E}^*}$ takes value zero. This would contradict the compactness of $B_{\mathcal{E}^*}$'s proper domain.

By (3) we know that productive aligned equilibria exist if Assumption 2 holds. To see the necessity of the assumption, suppose for a contradiction that it doesn't hold and yet some productive aligned equilibrium exists. Let v be the agent's continuation value at some

[39] Since a weighted average of two nonnegative numbers can be zero only if both numbers are zero.

[40] Let $v_0 = \hat{v}$ and $v_{n+1} = v_n + \frac{1-\delta}{\delta}(v_n - \omega) \geq \hat{v} + n(\hat{v} - \omega)$.

on-path history at which the principal delegates. By (2), we know $v \leq \omega$. As Assumption 2 fails, we then know $\delta v < \underline{\theta}$, contradicting agent IC. □

Optimal Equilibrium Now, focus on the frontier of the whole equilibrium set. Before proceeding to the full characterization, we establish a single crossing result: indiscriminate project adoption is initially used only for the highest agent values.

Lemma 10. *Fix $B \in \mathcal{B}$, and suppose $B^{-1}(0) = [0, \omega]$:*

1. *If $v > \omega$, then $T_0 B(v) < T_f B(v)$ (unless both are ∞).*

2. *There is a cutoff $\underline{v} \geq \omega$ such that*

$$\begin{cases} T_0 B(v) \leq T_1 B(v) & \text{if } v \in [\omega, \underline{v}); \\ T_0 B(v) \geq T_1 B(v) & \text{if } v > \underline{v}. \end{cases}$$

Proof. $B(\omega) = 0$, and B is convex. Therefore, B is strictly increasing above ω on its domain, so that[41] $\check{T}_f B(v) > \check{T}_0 B(v)$, confirming the first point.

Given v,

$$\frac{v - (1-\delta)[\theta_E - \underline{\theta}]}{\delta} - \frac{v - (1-\delta)\theta_E}{\delta} = \frac{1-\delta}{\delta}\underline{\theta}$$

is a nonnegative constant.

Since B is convex, it must be that the continuous function

$$v \mapsto \check{T}_0 B(v) - \check{T}_1 B(v) = \delta(1-h)\left[B\left(\frac{v - (1-\delta)[\theta_E - \underline{\theta}]}{\delta}\right) - B\left(\frac{v - (1-\delta)\theta_E}{\delta}\right)\right] - (1-h)$$

is increasing on its proper domain. The second point follows.[42] □

Our characterization of the equilibrium frontier B can now be stated.

Lemma 11 (Equilibrium Frontier)**.**
Suppose Assumptions 1 and 2 hold. Let $B := B_{\mathcal{E}^*}$ and $\bar{v} := \max dom(B)$.

1. *$\bar{v} \geq \omega$, and $B(v) = 0$ for $v \in [0, \omega]$.*

[41]The relationship is as shown if $\check{T}_0 B(v) < \infty$. Otherwise, $T_f B(v) = T_0 B(v) = \infty$.
[42]Because wherever $\check{T}_i B(v) \geq \check{T}_j B(v)$, we have $T_i B(v) \geq T_j B(v)$ as well.

2. *If* $\bar{v} > \omega$, *then*

$$\begin{aligned} B(v) &= T_0 B(v) \text{ for } v \in \left[\omega,\ \delta\bar{v} + (1-\delta)(\theta_E - \underline{\theta})\right]; \\ B(v) &\text{ is affine in } v \text{ for } v \in \left[\delta\bar{v} + (1-\delta)(\theta_E - \underline{\theta}),\ \bar{v}\right]; \\ B(\bar{v}) &= T_1 B(\bar{v}). \end{aligned}$$

3. *If* $\bar{v} > \omega$, *then* $\pi(\bar{v}, B(\bar{v})) = 0$.

Proof. The first point follows directly from Lemma 9. Now suppose $\bar{v} > \omega$.
Let $\underline{v} := \delta\bar{v} + (1-\delta)(\theta_E - \underline{\theta})$. Any $v > \underline{v}$ has $\frac{v-(1-\delta)[\theta_E-\underline{\theta}]}{\delta} > \frac{\underline{v}-(1-\delta)[\theta_E-\underline{\theta}]}{\delta} = \bar{v}$, so that $B\left(\frac{v-(1-\delta)[\theta_E-\underline{\theta}]}{\delta}\right) = \infty$, and therefore (appealing to Lemma 7) $T_0 B(v) = \infty$. Therefore, the cutoff defined in Lemma 10 is $\leq \underline{v}$.

Since $T_0 B, T_1 B$ are both convex, there exist some $v_0, v_1 \in [\omega, \bar{v}]$ such that $v_0 \leq v_1, \underline{v}$, and:

$$\begin{aligned} B(v) &= 0 \text{ for } v \in [0, \omega]; \\ B(v) &= T_0 B(v) \text{ for } v \in [\omega, v_0]; \\ B(v) &\text{ is affine in } v \text{ for } v \in [v_0, v_1]; \\ B(v) &= T_1 B(v) \text{ for } v \in [v_1, \bar{v}]. \end{aligned}$$

Let $m > 0$ denote the left-sided derivative of B at v_1 (which is simply the slope of B on (v_0, v_1) if $v_0 \neq v_1$).

Let $[v_0, v_1]$ be maximal (w.r.t. set inclusion) such that the above decomposition is still correct.

Notice then that $v_1 = \bar{v}$. Indeed:

- If $\frac{v_1-(1-\delta)\theta_E}{\delta} \geq v_0$, then (appealing to Lemma 7) the right-side derivative $T_1 B' = m$ in some neighborhood of v_1 in $[v_1, \bar{v}]$. By convexity of B, this would imply $B(v) = T_1 B(v) = B(v_1) + m(v - v_1)$ in said neighborhood. Then, by maximality of v_1, it must be that $v_1 = \bar{v}$.

- If $\frac{v_1-(1-\delta)\theta_E}{\delta} < v_0$, then minimality of v_0 implies (again using Lemma 7) that the right-sided derivative $B'(v_1) = TB'(v_1) < m$ if $v_1 < \bar{v}$. As B is convex, this can't happen. Therefore, $v_1 = \bar{v}$.

Finally, we need to show that $v_0 = \underline{v}$. Now, by minimality of v_0, it must be that for any $v \in [0, v_0)$, the right-side derivative $B'(v) < m$. If $v_0 < \underline{v}$ (so that $\frac{v_0-(1-\delta)[\theta_E-\underline{\theta}]}{\delta} < \underline{v}$), then

Lemma 7 gives us

$$\begin{aligned}
m &= B'(v_0) \\
&\leq T_0 B'(v_0) \\
&= hB'\left(\frac{v_0 - (1-\delta)\theta_E}{\delta}\right) + (1-h)B'\left(\frac{v_0 - (1-\delta)[\theta_E - \underline{\theta}]}{\delta}\right) \\
&= hB'\left(\frac{v_0 - (1-\delta)\theta_E}{\delta}\right) + (1-h)m \\
&< m,
\end{aligned}$$

a contradiction. Therefore $v_0 = \underline{v}$, and the second point of the theorem follows.

For the last point, assume $\bar{v} > \omega$ and yields strictly positive profits. Then, for sufficiently small $\gamma > 0$, the function $B^\gamma \in \mathcal{B}$ given by $B^\gamma(v) = \begin{cases} B(v) & \text{if } v \in [0, \bar{v}] \\ T_1 B(v) & \text{if } v \in [\bar{v}, \bar{v} + \gamma] \end{cases}$ is self-generating, contradicting the fact that $\mathcal{E}^*$ is the largest self-generating set. □

To better understand efficient equilibrium behavior, the following lemma is useful.

Lemma 12. *Let $\bar{v}, B$ be as in Lemma 11. If $\bar{v} > \omega$, then there are constants $\alpha, \beta > 0$ $B(\omega + \epsilon) = \alpha\epsilon^{1+\beta}$ for sufficiently small $\epsilon > 0$.*

Proof. For $\epsilon \in (0, \bar{v} - \omega)$, define $Q(\epsilon) := \frac{B(\omega+\epsilon) - B(\omega)}{\epsilon} = \frac{1}{\epsilon} B(\omega + \epsilon)$. For small enough $\epsilon > 0$,

$$Q(\epsilon) = \delta\left[(1-h)B\left(\frac{(\omega+\epsilon) - (1-\delta)\omega}{\delta}\right) + h0\right] = (1-h)Q\left(\frac{\epsilon}{\delta}\right).$$

For low enough $t \in (-\infty, \log(\bar{v} - \omega))$, let $q(t) := \log Q(e^t)$. Then q is a continuous function, and $q(t - \log\delta) = q(t) - \log(1-h)$ for low enough $t \in \mathbb{R}$. Therefore, far enough to the left of its domain, q is affine with slope $\beta := \frac{\log(1-h)}{\log\delta} > 0$. That is, there is a constant α such that $q(t) = \log\alpha + \beta t$ for low enough t. Therefore, for small enough $\epsilon > 0$,

$$Q(\epsilon) = e^{q(\log\epsilon)} = e^{\log\alpha + \beta\log\epsilon} = \alpha\epsilon^\beta \implies B(\omega + \epsilon) = \alpha\epsilon^{1+\beta}.$$

Lastly, nonnegativity of B and Proposition 1 tell us $\alpha > 0$. □

Lastly, the next lemma records a sufficient condition for existence of some equilibrium with an initial bad project.

Lemma 13. *Suppose Assumptions 1, 2, and 3 all hold, and let $\bar{v}, B$ be as defined in Lemma 11. Then $\bar{v} > \omega$.*

Proof. By Lemma 9, we know $B(\omega) = 0$, implying (in the notation of Lemma 11)

$$\check{T}_1 B((1-\delta)\theta_E + \delta\omega) = (1-\delta)(1-h) + B(\omega) = (1-\delta)(1-h).$$

Assumption 3 then tells us that $\pi\left((1-\delta)\theta_E + \delta\omega,\ (1-\delta)(1-h)\right) \geq 0$, so that $T_1 B((1-\delta)\theta_E + \delta\omega) = \check{T}_1 B((1-\delta)\theta_E + \delta\omega) < \infty$. Therefore, $\bar{v} \geq (1-\delta)\theta_E + \delta\omega > \omega$. □

9.5 Proof of Theorem 2

We now complete the proof of Theorem 2, under the hypothesis that Assumptions 1, 2, and 3 all hold.

Proof. Let $\bar{v}, B$ be the highest agent value and the efficient frontier function, as in the statement of Lemma 11. In view of Lemmata 11 and 13, all that remains is to show that the function B is strictly convex on $\left[\omega,\ \delta\bar{v} + (\theta_E - \underline{\theta})\right]$, and that the equilibrium value set is in fact $\{(v,b) \in [0,\bar{v}] \times \mathbb{R} :\ B(v) \leq b \leq \xi v\} = \{(v,b) \in [0,\bar{v}] \times \mathbb{R} :\ b \geq B(v),\ \pi(v,b) \geq 0\}$.

It suffices to note the following:

- There is some $v_0 \in (\omega,\ \delta\bar{v} + (1-\delta)\omega)$ such that B is strictly convex on (ω, v_0)
- For every $v \in (\omega,\ \delta\bar{v} + (1-\delta)\omega)$, there is some $\tilde{v}(v) \in (\omega, v)$ such that B is strictly convex on $[\omega, v]$ if it is strictly convex on $[\omega, \tilde{v}(v)]$.

Indeed, these together imply that $\max\{v \in [\omega,\ \ \delta\bar{v} + (1-\delta)\omega] :\ B \text{ is strictly convex on } [\omega, v]\}$ (which must exist by continuity) is $\delta\bar{v} + (1-\delta)\omega$.

Lemma 12 guarantees the former condition, and the latter condition comes from the functional form of B, letting $\tilde{v}(v) := \frac{v-\omega-(1-\delta)\theta}{\delta}$. This guarantees strict convexity.

Finally, notice that, by the definition of $B = B_{\mathcal{E}^*}$ and principal participation,

$$\text{co}\{(v, B(v)) :\ v \in [0,\bar{v}]\} \subseteq \mathcal{E}^* \subseteq \{(v,b) \in [0,\bar{v}] \times \mathbb{R} :\ b \geq B(v),\ \pi(v,b) \geq 0\}.$$

To show that the three sets are equal, consider any (v,b) in the last set with $b > B(v)$. As π is linear, we know $\pi\left(\bar{v}, \frac{\bar{v}}{v}b\right) \geq 0$. As $\pi(\bar{v}, B(\bar{v})) = 0$, it must be that $\frac{\bar{v}}{v}b \leq B(\bar{v})$. Therefore, the intermediate value theorem tells us there is some $v^1 \in (v, \bar{v}]$ such that $B(v^1) = \frac{v^1}{v}b$. Therefore, $(v,b) \in \text{co}\left\{(0,0),\ \left(v^1, \frac{v^1}{v}b\right)\right\} \subseteq \text{co}\{(v, B(v)) :\ v \in [0,\bar{v}]\}$. □

10 APPENDIX: Delayed Differential Equation

Taking a change of variables, from agent value v to account balance $x = \frac{v-\omega}{\underline{\theta}}$, the following system of equations describes the frontier of the equilibrium value set.

$$
\begin{aligned}
(1+\eta)b(x) &= \eta b(x-1) + xb'(x) \text{ for } x > 0 \\
b(x) &= 0 \text{ for } x \leq 0
\end{aligned}
$$

Theorem 4. *Consider the above system of equations. For any $\alpha \in \mathbb{R}$, there is a unique solution $b^{(\alpha)}$ to the above system with $b^{(\alpha)}(1) = \alpha$. Moreover $b^{(\alpha)} = \alpha b^{(1)}$.*
Letting $b = b^{(1)}$:

1. *$b(x) = x^{1+\eta}$ for $x \in [0, 1]$.*
2. *b is twice-differentiable on $(0, \infty)$ and globally C^1 on $\mathbb{R}$.*
3. *b is strictly convex on $(0, \infty)$ and globally convex on $\mathbb{R}$. In particular, b is unbounded.*
4. *b is strictly increasing and strictly log-concave on $(0, \infty)$.*

Proof. First consider the same equation on a smaller domain,

$$(1+\eta)b(x) = xb'(x) \text{ for } x \in (0, 1].$$

As is standard, the full family of solutions is $\{b^{(\alpha,1)}\}_{\alpha\in\mathbb{R}}$. where $b^{(\alpha,1)}(x) = \alpha x^{1+\eta}$ for $x \in (0, 1]$.

Now, given a particular partial solution $b : (-\infty, z] \to \mathbb{R}$ up to $z > 0$, there is a unique solution to the first-order linear differential equation $\hat{b} : [z, z+1] \to \mathbb{R}$ given by

$$\hat{b}'(x) = \frac{1+\eta}{x}\hat{b}(x) - \frac{\eta}{x}b(x-1).$$

Proceeding recursively, there is a unique solution to the given system of equations for each α. Moreover, since multiplying any solution by a constant yields another solution, uniqueness implies $b^{(\alpha)} = \alpha b^{(1)}$. Now let $b := b^{(1)}$.

We have shown that $b(x) = x^{1+\eta}$ for $x \in [0, 1]$, from which it follows readily that b is $C^{1+\lfloor\eta\rfloor}$ on $(-\infty, 1)$, Given $x > 0$, for small ϵ,

$$\begin{aligned}
(x+\epsilon)\frac{b'(x+\epsilon)-b'(x)}{\epsilon} &= \frac{1}{\epsilon}(x+\epsilon)b'(x+\epsilon) - \frac{1}{\epsilon}xb'(x) - b'(x) \\
&= \frac{1}{\epsilon}\Big[(1+\eta)b(x+\epsilon) - \eta b(x+\epsilon-1)\Big] - \frac{1}{\epsilon}\Big[(1+\eta)b(x) - \eta b(x-1)\Big] - b'(x) \\
&= \eta\left[\frac{b(x+\epsilon)-b(x)}{\epsilon} - \frac{b(x-1+\epsilon)-b(x-1)}{\epsilon}\right] + \left[\frac{b(x+\epsilon)-b(x)}{\epsilon} - b'(x)\right] \\
&\xrightarrow{\epsilon\to 0} \eta[b'(x) - b'(x-1)] + 0.
\end{aligned}$$

So b is twice differentiable at $x > 0$ with $b''(x) = \dfrac{\eta}{x}[b'(x) - b'(x-1)]$.

Let $\bar{x} := \sup\{x > 0 : \ b'|_{(0,x]} \text{ is strictly increasing}\}$. We know $\bar{x} \geq 1$, from our explicit solution of b up to 1. If $\bar{x}$ is finite, then $b'(\bar{x}) > b'(\bar{x}-1)$. But then $b''(\bar{x}) = \frac{\eta}{x}[b'(x) - b'(x-1)] > 0$, so that b' is strictly increasing in some neighborhood of $\bar{x}$, contradicting the maximality of $\bar{x}$. So $\bar{x} = \infty$, and our convexity result obtains. From that and $b'(0) = 0$, it is immediate that b is strictly increasing on $(0, \infty)$.

Lastly, let $f := \log b|_{(0,\infty)}$. Then $f(x) = (1+\eta)\log x$ for $x \in (0, 1]$, and for $x \in (1, \infty)$,

$$\begin{aligned}
(1+\eta)e^{f(x)} &= \eta e^{f(x-1)} + xe^{f(x)}f'(x), \\
\implies (1+\eta) &= \eta e^{f(x-1)-f(x)} + xf'(x). \\
\implies 0 &= \eta e^{f(x-1)-f(x)}[f'(x-1) - f'(x)] + f'(x) + f''(x) \\
\implies -f''(x) &= \eta e^{f(x-1)-f(x)}[f'(x-1) - f'(x)] + f'(x) \\
&\geq \eta e^{f(x-1)-f(x)}[f'(x-1) - f'(x)], \text{ since } f = \log b \text{ is increasing.}
\end{aligned}$$

The same contagion argument will work again. If f has been strictly concave so far, then $f'(x) < f'(x-1)$, in which case $-f''(x) > 0$ and f will continue to be concave. Since we know $f|_{(0,1]}$ is strictly concave, it follows that f is globally such. □

The first point of the following proposition shows that the economically intuitive boundary condition of our DDE uniquely pins down the solution b for any given account cap. The second point shows that as the account cap increases, so does the number of bad projects (in expected discounted terms) anticipated at the cap.

Proposition 4. *For any $\bar{x} > 0$*

1. *There is a unique $\alpha = \alpha(\bar{x}) > 0$ such that $b^{(\alpha)}(\bar{x}) = 1 + b^{(\alpha)}(\bar{x} - 1)$.*
2. *$b^{\alpha(\bar{x})}(\bar{x})$ is increasing in $\bar{x}$.*

Proof. The first part is immediate, with $\alpha = \dfrac{1}{b^{(1)}(\bar{x}) - b^{(1)}(\bar{x}-1)}$.

For the second part, notice that $\dfrac{b(\bar{x})}{b(\bar{x}-1)}$ is decreasing in $\bar{x}$ because b is log-concave. Then,

$$b^{\alpha(\bar{x})}(\bar{x}) = \alpha(\bar{x})b(\bar{x}) = \frac{b(\bar{x})}{b(\bar{x}) - b(\bar{x}-1)} = \frac{1}{1 - \frac{b(\bar{x}-1)}{b(\bar{x})}}$$

is increasing in $\bar{x}$. □

11 APPENDIX: Comparative Statics

In this section, we prove Proposition 3.

For any parameters $\eta, \bar{\theta}, \underline{\theta}, c$ satisfying Assumptions 1 and 2, and for any balance and bad projects x, b satisfying $x \geq b > 0$, define the associated profit

$$\begin{aligned}
\hat{\pi}(x, b|\eta, \bar{\theta}, \underline{\theta}, c) &:= \eta(\bar{\theta} - \underline{\theta}) + \underline{\theta}x - c\left[\frac{\eta(\bar{\theta} - \underline{\theta}) + \underline{\theta}x - \underline{\theta}b}{\bar{\theta}} + b\right] \\
&= \left(1 - \frac{c}{\bar{\theta}}\right)[\eta(\bar{\theta} - \underline{\theta}) + \underline{\theta}x] - c\left(1 - \frac{\underline{\theta}}{\bar{\theta}}\right)b \\
&= (\bar{\theta} - c)\eta + \left(1 - \frac{c}{\bar{\theta}}\right)\underline{\theta}(x - \eta) - c\left(1 - \frac{\underline{\theta}}{\bar{\theta}}\right)b.
\end{aligned}$$

For reference, we compute the following derivatives of profit:

$$\begin{aligned}
\frac{\partial\hat{\pi}}{\partial\bar{\theta}} &= \eta - c\underline{\theta}\frac{-1}{\bar{\theta}^2}[x - b - \eta] = \left(1 - \frac{c\underline{\theta}}{\bar{\theta}^2}\right)\eta + \frac{c\underline{\theta}}{\bar{\theta}^2}(x - b) > 0. \\
\frac{\partial\hat{\pi}}{\partial\underline{\theta}} &= \left(1 - \frac{c}{\bar{\theta}}\right)(x - \eta) + c\frac{1}{\bar{\theta}}b, \text{ which implies} \\
(\bar{\theta} - \underline{\theta})\frac{\partial\hat{\pi}}{\partial\underline{\theta}} + \hat{\pi} &= (\bar{\theta} - c)\eta + \left(1 - \frac{c}{\bar{\theta}}\right)\bar{\theta}(x - \eta) - 0b = (\bar{\theta} - c)x > 0. \\
\frac{\partial\hat{\pi}}{\partial c} &= -\frac{1}{\bar{\theta}}\left[\eta(\bar{\theta} - \underline{\theta}) + \underline{\theta}x + (\bar{\theta} - \underline{\theta})b\right] < 0.
\end{aligned}$$

Fix parameters $(\bar{\theta}, \underline{\theta}, c)$, and let $\bar{x}^*$ be as delivered in Theorem 3. We first show that slightly raising either of $\bar{\theta}, \underline{\theta}$ or slightly lowering c weakly raises the cap, strictly if $\bar{x}^* > 0$.

- If $\bar{x}^* = 0$, there is nothing to show, so assume $\bar{x}^* > 0$ henceforth. Notice that the expected discounted number of bad projects when at the cap depends only on η and the size of the cap. By Theorem 3, $\hat{\pi}(\bar{x}^*, b|\eta, \bar{\theta}, \underline{\theta}, c) = 0$.

- Consider slightly raising $\bar{\theta}$ or $\underline{\theta}$, or slightly lowering c. By the above derivative computations, the profit of the DCB contract with cap $\bar{x}^*$ is strictly positive.

- For any of the above considered changes, the DCB contract with cap $\bar{x}^*$ has strictly positive profits. Appealing to continuity, a slightly higher cap still yields positive profits under the new parameters, and is therefore consistent with equilibrium by Proposition 2. Then, appealing to Theorem 3 again, the cap associated with the new parameters is strictly higher than $\bar{x}^*$.

Now we consider comparative statics in the profit-maximizing initial account balance. We will fix parameters $(\bar{\theta}, \underline{\theta}, c)$, and consider raising either of $\bar{\theta}, \underline{\theta}$ or lowering c.

- Again, if $\bar{x}^* = 0$ at the original parameters, there's nothing to check, so assume $\bar{x}^* > 0$.

- Let $\check{b}$ be some solution to the DDE in Section 10, so that the expected discounted number of bad projects at a given cap $\bar{x}$ and balance x is $b(x|\bar{x}) = \dfrac{\check{b}(x)}{\check{b}(\bar{x}) - \check{b}(\bar{x}-1)}$. Because $\check{b}$ is strictly increasing and strictly convex (by work in Section 10), we know that $\frac{\partial}{\partial x} b(x|\bar{x})$ is strictly decreasing in $\bar{x}$. Therefore, by our comparative statics result for the cap, the parameter change results in a global strict decrease of $\frac{\partial}{\partial x} b(x|\bar{x}^*)$.

- By the form of $\hat{\pi}$ and by convexity of $b(\cdot|\bar{x})$, the unique optimal initial balance is the balance at which $\frac{\partial}{\partial x} b(x|\bar{x})$ is equal to

$$\xi = \frac{\underline{\theta}\left(1 - \frac{c}{\bar{\theta}}\right)}{c\left(1 - \frac{\underline{\theta}}{\bar{\theta}}\right)} = \frac{\underline{\theta}(\bar{\theta} - c)}{c(\bar{\theta} - \underline{\theta})},$$

  which increases with the parameter change.

- As ξ increases and $\frac{\partial}{\partial x} b(x|\bar{x}^*)$ decreases (at each x) with the parameter change, the optimal balance x^* must increase (given convexity) to satisfy the first-order condition $\frac{\partial}{\partial x}\Big|_{x=x^*} b(x|\bar{x}) = \xi$.