Repeated Delegation

by *Elliot Lipnowski & João Ramos*[*]

New York University

February 7, 2015

Abstract

We consider an ongoing relationship of delegated decision making. A principal, facing a stream of projects to potentially finance, must rely on an agent to assess the returns of different opportunities. As the cost of initiating a project is borne by the principal alone, the players disagree about which projects are worth financing. That the principal cannot commit limits the rewards she can credibly offer the agent for his fiscal restraint. Even so, we show that the principal can credibly—and indeed, should—employ the promise of some bad projects (future lenience) to incentivize the agent. We characterize the optimal contract, termed *Dynamic Capital Budgeting*, which consists of two distinct regimes. In the first regime, Capped Budgeting, the principal allocates an expense account (populated by "funny money") to the agent and fully delegates project choice, funded from the account; the account grows at the interest rate so long as its balance stays below a given cap. Only at the cap, where the account can grow no further, the agent is inconsiderate of the principal's interests. After enough projects have been initiated, a Controlled Budgeting regime begins, and the agent loses his autonomy forever.

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

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

This work has benefited from discussions with Heski Bar-Isaac, V. Bhaskar, Adam Brandenburger, Sylvain Chassang, Joyee Deb, Ross Doppelt, Ignacio Esponda, Eduardo Faingold, Vijay Krishna, Laurent Mathevet, David Pearce, Debraj Ray, Alejandro Rivera, Ariel Rubinstein, Evan Sadler, Maher Said, Tomasz Sadzik, Ennio Stacchetti, Bruno Strulovici, Laura Veldkamp, Larry White, Alex Wolitzky, Chris Woolnough, Huseyin Yildirim, Sevgi Yuksel, and seminar participants at NYU and EconCon (Princeton). 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. 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 Defense funds specialized equipment for each of its units, but must rely on those on the ground to assess their need for it. Relationships organized through delegation are often ongoing, consisting of many distinct decisions to be made over time; our focus is how this feature affects the relationship.

Beyond the absence 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.

Even for decisions that are in principle separate—in which today's course of action has no bearing on tomorrow's prospects—the controlling party can connect them, as a means to discipline the informed party now. If the university restricts the physics department to ten hires per decade, this might persuade the physics department to be discerning in the present; hiring a mediocre physicist would crowd out a good one. By employing a *budgeting* rule, the controlling party imposes a cost on the agent for excessive spending, better aligning their interests.

We study an infinitely repeated game between a principal (♀) with full authority over a decision to be made in an uncertain world; she relies on an agent (♂) to assess the state. Each period, the principal must choose whether or not to initiate a project, which may be good (i.e. high value enough to offset its cost) or bad. The principal herself is ignorant of the current project's quality, but the agent knows 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 prefers to fund only good projects, while the agent always prefers to invest in a 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 to fund it or not. The park would surely benefit the county, but the state must weigh this benefit against the money's opportunity cost. To asses this trade-off, 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. If the county were never selective in its proposals, the state would never want to fund the park. The agent's private information is transient: project types are independent across time, and a given project only affects within-period payoffs.[1]

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's built. If, further, 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 the agent can expect from the principal in the future. The state government can't credibly reward a county's fiscal restraint today by promising carte blanche in the future.

The present conflict of interests would be resolved if the principal could sell permanent control to the agent.[2] In keeping with our leading applications, we focus on the repeated interaction without monetary transactions.[3] The Department of Defense seems unlikely to ask soldiers to pay for their own body armor.

Our first result is an efficiency bound on any delegation rule that involves only good projects being taken. If no bad projects are initiated, the principal and the agent have aligned interests, both preferring more projects. The best such delegation rule, the **Aligned Optimal Budget**, has a very simple form. The principal delegates to the agent until the agent adopts a project, but follows any project with a temporary freeze. That is, no more projects are allowed for the next $\bar{\tau}$ units of time. Then the same contract starts over. The optimal $\bar{\tau}$ will be the shortest freeze time sufficiently severe to keep the agent from taking bad projects. In the context of a university, the physics department can freely search for a candidate, but any hire is followed by a temporary hiring freeze. During the freeze, many qualified candidates may be available, but the department is forbidden from hiring them. We show that the resulting inefficiency remains even if the parties are arbitrarily patient.

Our main result is a full characterization of the optimal intertemporal delegation rule. The optimal contract, the **Dynamic Capital Budget**, comprises two distinct regimes. At any time, the parties engage in either Capped Budgeting or Controlled Budgeting.

---

[1] That is, we abstract from the intrinsic dynamic consequences of adopting a project, e.g. affecting the remaining pool of potential projects, or affecting the needs/preferences of the agent going forward. In this sense, we isolate the dynamic delegation problem.

[2] This standard solution is sometimes called "selling the firm". For instance, see (Mas-Colell, Whinston, and Green, 1995, p. 482).

[3] This assumption is stronger than needed. As long as the agent cannot make transfers to the principal, our main results are qualitatively unchanged.

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 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 doesn't grow any larger (even if he waits). Not being rewarded for fiscal restraint, the agent immediately initiates a project, and again his balance declines by 1.

If the agent overspends, a **Controlled Budget** regime begins: the principal first temporarily freezes, a larger overdraft being met with a longer freeze. The parties then revert to the Aligned Optimal Budget. It's worth noting that the players are certain to eventually enter this regime. Once there, the Controlled Budget regime is absorbing.

This result is surprising in two distinct ways. First, our disadvantaged principal, stripped of all her usual tools, can nonetheless leverage the agent's information with some success. Second, the optimal equilibrium exhibits remarkably rich dynamics for such a simple, stationary model.

Prima facie, the only incentivizing instrument available to the principal is mutual "money burning", in the form of (temporarily) freezing. A freeze on equipment acquisition by the Department of Defense can be a useful threat, inducing frugal decisions now, but it comes at a cost: its own forces will sometimes be unequipped even in times of need. This force, with the cost it entails, single-handedly disciplines the agent under Controlled Budgeting. However, the principal has an additional tool: the expectation of future lenience can serve as a reward for the agent today. To induce frugal decisions now, the Department of Defense may promise more budgetary freedom in the future. The Capped Budget makes use of both this reward and above punishment—*carrot* and *stick*. A high account balance entails the promise of future permissiveness from the principal, while a low account balance entails an imminent threat of Controlled Budgeting. When the budget is below the cap, the principal rewards the agent for his diligence with the account's interest accrual. As long as the promise is credible—i.e. the principal would rather fulfill her contract than unilaterally freeze the relationship—the reward will be credible too. At the cap itself, the principal cannot credibly promise further lenience, and good behavior by the agent would go unpaid; accordingly the agent takes a project immediately. If the unit has shown enough fiscal restraint, the Department of Defense purchases new equipment, independent of its need, to reward the unit.

The optimal contract yields clear dynamics for the delegation relationship. Both regimes reflect a productive relationship, but each is of a distinct character. Capped Budgeting is highly productive but low-yield:[4] every good project is being adopted, but some bad projects are as well. Controlled Budgeting is high-yield but less productive: only good projects are adopted, but some good opportunities go unrealized. In this sense, as the Capped Budget regime is transient, the relationship naturally drifts toward conservatism. The principal's payoff comparisons among regimes are ambiguous: at lower budget balances, Capped Budgeting dominates Controlled Budgeting, while for larger balances the relationship reverses.

The remainder of the paper is structured as follows. In the following pages, we discuss the related literature. Section 2 presents the model. In Section 3, we introduce convenient language for discussing players' incentives in our model. In Section 4, 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 5 studies contracts based on an annual budget, which we show they can be better or worse than aligned equilibria. Section 6 presents our optimal contract. In Section 7, we discuss some possible extensions of our model. Final remarks follow in Section 8.

Related Literature

This paper belongs to a rich literature on delegated decision-making,[5] initiated by Holmström (1984), wherein a principal faces a tradeoff between leveraging an agent's private information and shielding herself from his conflicting interests. The key issue is how much freedom the principal should give the agent; the more aligned are their preferences, the more discretion he should have. Armstrong and Vickers (2010) find that the principal optimally excludes some ex-post favorable options, in order to provide better incentives to the agent ex-ante, while Ambrus and Egorov (2013) highlight the value created by money burning as a means to alleviate incentive constraints. These insights apply to our model, in which indirect money burning—harmful to both players—is used to provide incentives.

Our paper contributes to the recently active field of dynamic delegation. Malenko (2013) characterizes the optimal contract for a principal who delegates investment choices and has a costly state verification technology, monetary transfers, and commitment power: a capital expense account with fluctuating interest rate. Guo (2014) focuses on delegation of a dynamic decision problem to an agent with non-transient private information. Alonso and Matouschek (2007) indicate how dynamic threats can partially bridge the gap between the cheap-talk and delegation models. In contemporaneous work, Guo and Hörner (2014) study optimal dynamic mechanisms without money in a world of partially persistent valuations,

[4]By "productive", we mean a lot of value is being delivered to the agent. By "low-yield", we mean less such value is delivered per unit of principal cost.

[5]For instance, see Frankel (2014) and the thorough review therein of the delegation literature.

in which the principal has commitment power. The principal's ability to commit generates different incentive dynamics: in contrast to our model, the agent may receive his first-best outcome in the long-run.

Our model speaks to the literature on relational contracting, as in Pearce and Stacchetti (1998), Levin (2003), and Malcomson (2010). The relational contract literature focuses on relationships in which formal contracting is impossible, and all incentives—and the credibility of promises that provide those incentives—are anchored to the future value of the relationship. In particular, Li and Matouschek (2013) focus on the case in which the principal's opportunity cost of promise keeping is private information. In both their model and ours, the inability to formally contract leads a stationary problem to be met with a non-stationary relationship. In theirs, the relationship is cyclical, with every punishment being strictly temporary. In ours, the relationship temporarily cycles, before drifting toward conservatism.

Our results add to the literature on relationship building under private information. One strand of the literature concerns itself with building and maintenance of partnerships, for instance Möbius (2001), Hauser and Hopenhayn (2008), and Espino, Kozlowski, and Sanchez (2013). Möbius (2001) constructs a model in which players privately observe opportunities to do favors for one another at personal cost; Hauser and Hopenhayn (2008) indicate that the relationship can benefit from varying incentives based on both action and inaction. In a related strand of the literature, Chassang (2010) and Li, Matoushek, and Powell (2014) focus on the relationship between a firm and its employee, whose private information can generate persistent differences in performance across ex-ante identical firms. Li, Matoushek, and Powell (2014) generates similar history dependence to Guo and Hörner (2014), in a repeated trust game setting. Owing to the firm's inability to interpret its employee's actions, the realization of random early outcomes can have long-lasting consequences. The publicly observed arrival of a new technology further highlights the potential extent of past burdens: the relationship's past may preclude the firm from adapting to dramatic changes in its environment. A final strand of this literature regards dynamic corporate finance, as in Clementi and Hopenhayn (2006) or Biais et al. (2010). In Biais et al. (2010), 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 ours exhibit similar dynamics: our "funny money" balance takes the role of real sunk investment.

Lastly, there is a deep connection between the present work and the literature on linked decisions. Casella (2005) and Jackson and Sonnenschein (2007) consider a settings in which, given a large number of physically independent decisions, the ability to connect them across time helps align incentives. Frankel (2011) and Frankel (2013) consider environments in which a principal with commitment power optimally employs a budgeting rule to discipline the agent.

2 The Model

We consider an infinite horizon two-player ($\mathcal{P}$rincipal & $\mathcal{A}$gent) game in discrete time. Each period, the principal chooses whether or not to delegate the project adoption choice to the agent. Conditional on delegation, the agent (privately) observes which type of project is available and (publicly) decides whether or not to adopt it. Any project of type θ generates an agent payoff of θ and a cost of c, and thus a principal payoff of $\theta - c$, at the time of the project's adoption. Notice that the cost is independent of 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 payoff, in addition to the cost that she alone bears. While the university's president can't 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 prefers those that benefit the community. Given this altruistic motive, the principal cares about the value generated by a project, a value she never observes.

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 only cares 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 $t \in \mathbb{Z}_+$ is θ_t and projects are adopted in periods $\mathcal{T} \subseteq \mathbb{Z}_+$, then the principal and agent get profit and revenue,

$$\Pi = \sum_{t \in \mathcal{T}} \delta^t (\theta_t - c) \text{ and } V = \sum_{t \in \mathcal{T}} \delta^t \theta_t, \text{ respectively.}$$

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

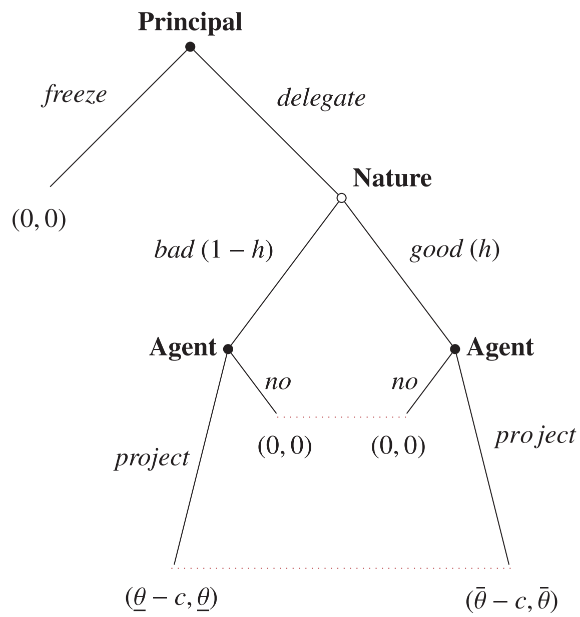

Figure 1: The principal observes agent choice but not project quality.

First, $\mathcal{P}$ publicly decides whether to *freeze* project adoption or to *delegate* it. If $\mathcal{P}$ freezes, nothing happens and both players accrue no payoffs. If $\mathcal{P}$ delegates, $\mathcal{A}$ privately observes which type of project is available and decides whether or not to initiate the best 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* $\theta_E := (1-h)\underline{\theta} + h\bar{\theta}$, *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 does still prefer good ones to bad ones. So they agree on which projects are best to adopt, but may disagree on whether a given project is worth taking ex-post. The condition $\theta_E < c$ (interpreted as an assumption that good projects are scarce) says the latter effect dominates, and the conflict of interests prevails even ex-ante: the principal prefers freeze to the average project. A good enough physicist is rare; the university only finds hiring worthwhile if it can rely on the department to separate the wheat from the chaff. If the players only interacted 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.

Equilibrium Values

Throughout the paper, *equilibrium* will be taken to mean perfect public equilibrium (PPE), in which the players respond only to the public history of actions and (for the agent) current project availability.

Definition. *A* ***perfect public equilibrium***[6] *of a discrete time repeated game is a sequential equilibrium, in semi-public strategies,*[7] *of the game augmented with public correlation at each history.*

For ease of bookkeeping, it's convenient to track equilibrium supported revenue v and bad projects b,[8] both in expected discounted terms. The vector (v, b) encodes both agent value v and principal profit

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

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

[6]See (Mailath and Samuelson, 2006, Definition 7.1.3) for a formal definition.

[7]i.e. $\mathcal{P}$ conditions only on public history, while $\mathcal{A}$ conditions only on the public history and the current project type.

[8]The more fundamental objects are (expected discounted) good projects $g = \sum_{t\in\mathcal{T}} \delta^t \mathbf{1}_{\{\theta_t=\bar{\theta}\}}$ and (expected discounted) bad projects $b = \sum_{t\in\mathcal{T}} \delta^t \mathbf{1}_{\{\theta_t=\underline{\theta}\}}$. Given those, we can readily compute the agent value/revenue as

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

and the principal value/profit as

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

Clearly, the pairs (g, b), (v, π), and (v, b) are all linear transformations of one another.

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

Throughout the paper we make extensive use of two simple observations about the set $\mathcal{E}^*$. First, notice that $\vec{0} \in \mathcal{E}^*$, 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. For any profile satisfying appropriate on-path incentive constraints, one can always find another profile with identical on-path behavior, but altered off-path to make the profile an equilibrium. With the lemma in hand, we don't specify off-path behavior in a given strategy profile, as we are chiefly interested in payoffs.

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

1. *All on-path incentive constraints are satisfied for the agent.*
2. *Following all on-path histories, the principal has nonnegative continuation profit.*

Then there is an equilibrium $\tilde{\sigma}$ which 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's prescription), play according to σ.
- Off-path (i.e. if $\mathcal{P}$ has ever deviated from σ's prescription), 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's also incentive-compatible for the principal: off-path because σ^{static} is, on-path because σ has nonnegative continuation profits while σ^{static} yields zero profit. □

3 Dynamic Incentives

While our results concern a discrete time repeated game, we find it expositionally convenient to present the intuition in continuous time. We present results for the case in which the players interact very frequently, but good projects remain scarce. A unit can find desirable equipment to request from the Department of Defense at any time; what's rare is the opportunity to buy equipment whose benefit offsets its cost. The cleanest economic intuition lies in this limiting case. Letting the time between decisions, together with the proportion of good projects, vanish enables us to present our main results heuristically in the language of calculus.

Given a period length $\Delta > 0$, we parametrize the discount factor $\delta = 1 - \Delta$, and the proportion of good projects $h = \eta\Delta$ for a fixed $\eta > 0$. At the limit, as $\Delta \to 0$, good projects arrive with Poisson rate η, and bad projects are always available. Finally, observe that in the limiting case $\theta_E = \underline{\theta}$. Even though we consider the limit as $\Delta \to 0$ and thus $\delta \to 1$, the present work is not a folk theorem analysis.[9]

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

- $\dot{v}_t$ [resp. $\dot{b}_t$], the rate of change of v [resp. b] conditional on no project adoption; and
- $\tilde{v}_t$ [resp. $\tilde{b}_t$], the continuation of v [resp. b] in case a project is undertaken.

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 as follows:

1. The principal makes a delegation choice[10] $d \in \{0, 1\}$, i.e. whether or not to let the agent pick up projects in the current period.
2. The agent chooses $\hat{\eta} \in [0, \eta]$ and $\hat{\lambda} \in [0, \infty]$, the instantaneous rates at which he currently initiates good and bad projects respectively, conditional on being allowed to.

By way of interpretation, $d_t, \hat{\eta}_t, \hat{\lambda}_t$ are the approximate choices on either $[t, t + \Delta)$ for small $\Delta > 0$ or until the next project, whichever comes first.[11] In particular, $d_t = 1$ doesn't mean that the agent has the opportunity to initiate an arbitrarily large number of bad projects without the principal's continued consent.

Appealing to self-generation arguments as in Abreu, Pearce, and Stacchetti (1990) and to Lemma 1, equilbrium is characterized by:

[9] A folk theorem analysis would entail taking $\delta \to 1$, for a fixed h. The distinction is analogous to that in Abreu, Milgrom, and Pearce (1991).

[10] As we show in the appendix, it's without loss of generality that the principal uses a pure strategy. The intuition is that, since everything $\mathcal{P}$ does is publicly observed, any mixing she does may as well be replaced with public mixing.

[11] Notice that, if the principal chooses $d_t = 1$ and the agent chooses $\hat{\lambda}_t = \infty$, then both players face new choices to make, still exactly at time t.

1. Promise keeping:[12]

$$\begin{aligned} v &= d\left[(\hat{\eta}\bar{\theta} + \hat{\lambda}\underline{\theta}) - (\hat{\eta} + \hat{\lambda})(v - \tilde{v})\right] + \dot{v} \\ &= d\hat{\eta}(\bar{\theta} + \tilde{v} - v) + d\hat{\lambda}(\underline{\theta} + \tilde{v} - v) + \dot{v} \\ b &= d\left[\hat{\lambda} - (\hat{\eta} + \hat{\lambda})(b - \tilde{b})\right] + \dot{b} \\ &= d\hat{\eta}(0 + \tilde{b} - b) + d\hat{\lambda}(1 + \tilde{b} - b) + \dot{b}. \end{aligned}$$

2. Agent incentive-compatibility:

$$v - \tilde{v} \begin{cases} \geq \underline{\theta} & \text{if } \hat{\lambda} < \infty \\ \leq \bar{\theta} & \text{if } \hat{\eta} > 0. \end{cases}$$

3. Principal participation:

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

Promise keeping decomposes the continuation outcomes (v, b) from any instant into what happens in each of three events—the agent finds and invests in a good project at that instant; the agent finds and invests in a bad project at that instant; no project is adopted—weighted by their instantaneous probabilities. The agent's incentive constraint states that, if the agent is willing to resist taking a project immediately ($\hat{\lambda} < \infty$) 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 ($\hat{\eta} > 0$), the same punishment $v - \tilde{v}$ cannot be too draconian. Finally, 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.

4 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 limit-

[12] When $d = 1$ and $\hat{\lambda} = \infty$, we replace the given equations with the limiting equations obtained from dividing through by $\hat{\lambda}$:

$$\begin{aligned} 0 &= \underline{\theta} - (v - \tilde{v}) \\ 0 &= 1 - (b - \tilde{b}). \end{aligned}$$

When $d = 0$ and $\hat{\lambda} = \infty$, we let $d\hat{\lambda} = 0$, so that the principal retains ultimate authority over project adoption.

less bad projects. In the present section, we ask whether this core tension can be resolved by allowing non-stationary equilibria. More precisely, are there productive aligned equilibria?

Definition. *An **aligned equilibrium** is an equilibrium in which there are no bad projects.*

A sensible first attempt is to delegate, but to punish the agent as much as possible as soon as he might have taken a bad project. Describe σ^∞ as follows: the principal allows exactly one project, after which she shuts down forever; the agent takes the first good project that comes along. Is σ^∞ an equilibrium? For this profile, before the first project,

$$d = 1, \;\; \hat{\eta} = \eta, \;\; \hat{\lambda} = 0, \;\; \tilde{v} = 0, \text{ and } \dot{v} = 0.$$

Therefore, promise keeping gives

$$v = \eta(\bar{\theta} + 0 - v) + 0 \implies v = \frac{\eta}{1+\eta}\bar{\theta} \;\; \in (0, \bar{\theta}).$$

Principal participation is immediate when there are no bad projects, so that we only need to check agent incentive-compatibility.

$$\text{Agent IC holds} \iff v - \tilde{v} \geq \underline{\theta} \iff \frac{\eta}{1+\eta}\bar{\theta} \geq \underline{\theta} \iff \eta(\bar{\theta} - \underline{\theta}) \geq \underline{\theta}.$$

Notation. *Let $\omega := \eta(\bar{\theta} - \underline{\theta})$ be the **marginal value of restraint**.*
Interestingly, the number ω appears throughout the analysis of the game.

Assumption 2.

$$\omega > \underline{\theta}.$$

Unless otherwise stated, we'll assume Assumption 2 holds throughout. In discrete time, Assumption 2 can equivalently be expressed as a lower bound on the discount factor δ. If the agent is sufficiently patient, the marginal value of restraint outweighs the myopic benefit of an immediate bad project.

A Stick with No Carrot

We've seen that the threat of shutdown is enough to incentivize picky project adoption. However, permanent shutdown destroys a lot of value, for both the principal and the agent. If the university only allows the physics department one hire for its entire existence, every good candidate is passed over thereafter, harming the university. It's natural to ask whether a less severe mutual punishment can provide the same incentives.

Given $\tau \in (0, \infty]$, describe the τ-freeze stationary contract σ^τ as follows:

1. The principal starts by delegating, and does so indefinitely if no projects are taken.

2. The agent takes no bad projects, and takes the first good project that arrives.

3. Any project is followed by a freeze of length τ, followed by restarting σ^τ.

We can interpret the τ-freeze contract as a simple budget rule. The agent is given a budget of one project by the principal. If $\mathcal{A}$ does not spend his budget, it rolls over to the next instant. If the budget is depleted, $\mathcal{P}$ replenishes it after a waiting period τ. Every hire by the department is followed by a two-year freeze, after which the university will again allow the department to hire.

As seen in the previous section, σ^∞ is an equilibrium. By continuity, σ^τ is an equilibrium for sufficiently high finite τ. Moreover, for $\tau \approx 0$, the contract σ^τ can't be an equilibrium. Indeed, in a delegation phase

$$v - \tilde{v} = (1 - e^{-\tau})v \leq (1 - e^{-\tau})\eta\bar{\theta} \xrightarrow{\tau \to 0} 0.$$

In particular, $(1 - e^{-\tau})\eta\bar{\theta} < \underline{\theta}$ for sufficiently small τ. If the punishment selecting a project is negligible, then the agent always strictly prefers to take the (bad) project in front of him.

It's clear that the revenue generated by σ^τ is decreasing in τ: less shutdown means fewer forgone opportunities for good projects, which means more revenue. What's less clear is how the punishment $v - \tilde{v}$ changes with τ. As we increase τ, the punishment is $v - \tilde{v} = (1 - e^{-\tau})v$, which increases as a fraction of total revenue. Thus its comparative statics are not obvious: increasing τ makes the punishment a bigger share of a smaller pie. As we'll see below, increasing τ always relaxes the agent's relevant IC constraint.

Let's analyze the value v and project punishment $v - \tilde{v}$ in the described contract. Promise keeping gives $v = \eta[\bar{\theta} + \tilde{v} - v] + 0 = \eta[\bar{\theta} - (1 - e^{-\tau})v]$, which implies

$$v = \frac{\eta}{1 + \eta(1 - e^{-\tau})}\bar{\theta} \text{ and } v - \tilde{v} = \frac{\eta(1 - e^{-\tau})}{1 + \eta(1 - e^{-\tau})}\bar{\theta}.$$

The revenue being decreasing in τ and the punishment being increasing in τ, the following proposition follows readily.

Proposition 1. *Consider* $\{\sigma^\tau\}_{\tau \in (0,\infty]}$ *as above.*

1. *There is a unique* $\bar{\tau} \in (0, \infty]$ *satisfying* $\dfrac{\eta(1 - e^{-\bar{\tau}})}{1 + \eta(1 - e^{-\bar{\tau}})}\bar{\theta} = \underline{\theta}$.

2. σ^τ *is an equilibrium if and only if* $\tau \geq \bar{\tau}$.

3. *Among all such* τ, *the choice* $\bar{\tau}$ *provides the highest revenue (and thus the highest profit).*

4. *The revenue of* $\sigma^{\bar{\tau}}$ *is* ω, *and so its profit is* $(1 - \frac{c}{\bar{\theta}})\omega$.

Proof. Everything is proven already except for the expression for $\bar{\tau}$ and the generated revenue. To compute $\bar{\tau}$,

$$\begin{aligned}
\frac{\eta(1-e^{-\tau})}{1+\eta(1-e^{-\tau})}\bar{\theta} = \underline{\theta} \quad &\iff \quad \eta(1-e^{-\tau})\bar{\theta} = [1+\eta(1-e^{-\tau})]\underline{\theta} \\
&\iff \quad (1-e^{-\tau})\omega = \underline{\theta} \\
&\iff \quad \tau = \log\frac{\omega}{\omega-\underline{\theta}}.
\end{aligned}$$

The associated revenue is then

$$\begin{aligned}
v &= \frac{\eta}{1+\eta(1-e^{-\bar{\tau}})}\bar{\theta} \\
&= \frac{\eta\bar{\theta}}{1+\eta\frac{\underline{\theta}}{\omega}} = \frac{\eta\bar{\theta}}{\omega+\eta\underline{\theta}}\omega \\
&= \omega.
\end{aligned}$$

□

This simple class of contracts illuminates the forces at play in our model. The principal wants good projects to be initiated, but she can't afford to give the agent free rein. If she wants to stop him from investing in bad projects, she must threaten him with mutual money burning. Subject to wielding a large enough stick to encourage good behavior, she efficiently leaves as little money on the table as possible.

One may be concerned that frequent shutdown leaves a lot of opportunities unrealized. Accordingly, it seems sensible to seek other plausible aligned equilibria in which less value is destroyed.

What if, instead of allowing one project followed by temporary freeze, $\mathcal{P}$ allows $K \in \mathbb{N}$ projects, before freezing for $\tau \in (0, \infty]$? The main lesson in Jackson and Sonnenschein (2007) is that budgetary rationing of multiple decisions can alleviate incentive misalignment, at a minor welfare cost. One might therefore hope that allowing K projects before punishing enables a more productive relationship, while still incentivizing the agent to avoid bad projects. As it turns out, this affords no real improvement. If the (now more distant) punishment for the first project is to be enough to stop the agent from cheating initially, it must be more severe, enough so to negate the would-be benefits of delayed closure. With similar computations to those in the previous section, it's straightforward that: the "K-project, τ-freeze" strategy profile is an equilibrium if and only if it delivers an initial value $\leq \omega$ to the agent. So the principal can allow more projects before punishing the agent (and herself), but if she is to still deter the agent from cheating, she has to make the punishment phase—which happens further in the future—longer for bigger K. This makes higher K redundant: no such equilibrium can outperform the "1-project, $\bar{\tau}$-freeze" equilibrium.

Constrained Optimality

The preceding analysis suggests a limit to how productive an aligned equilibrium can be. Indeed, with no bad projects, the principal has only one dimension—expected discounted good projects, or equivalently $\mathbb{E} \int e^{-t} \eta 1_{\{\text{open at time } t\}} \, \mathrm{d}t$—with which to provide incentives. Delaying a punishment, making it less likely, or making it less severe are all different physical instruments to alleviate the same money burning cost, but with the same adverse effect on agent incentives. The following theorem shows that the upper bound we've uncovered in specific classes of equilibria—the marginal value of restraint—is no coincidence.

Theorem 1 (Aligned Optimality). [13]

1. *Every aligned equilibrium generates equilibrium revenue $\leq \omega$.*
2. *There exist productive aligned equilibria if and only if*[14] *Assumption 2 holds.*
3. *Given Assumption 2: the $\bar{\tau}$-freeze contract, where $\bar{\tau} = \log \dfrac{\omega}{\omega - \underline{\theta}}$, yielding revenue ω, is optimal among all aligned equilibria.*

Proof. Following any history, in any aligned equilibrium,

$$\begin{aligned}
\dot{v} &= v - p\left[(\hat{\eta}\bar{\theta} + \hat{\lambda}\underline{\theta}) - (\hat{\eta} + \hat{\lambda})(v - \tilde{v})\right] \\
&= v - p\left[(\hat{\eta}\bar{\theta} + 0\underline{\theta}) - (\hat{\eta} + 0)(v - \tilde{v})\right] \\
&= v - p\hat{\eta}[\bar{\theta} - (v - \tilde{v})] \\
&\geq v - p\hat{\eta}[\bar{\theta} - \underline{\theta}] \text{ (since agent IC \& no bad projects } \implies v - \tilde{v} \geq \underline{\theta} \text{ if } p > 0) \\
&\geq v - \eta(\bar{\theta} - \underline{\theta}) \\
&= v - \omega.
\end{aligned}$$

[13] The discrete time counterpart is Proposition 2.

[14] We abstract from the knife-edge case $\omega = \underline{\theta}$.

So if $\epsilon := v_0 - \omega > 0$, then v indefinitely grows at rate $\dot{v} \geq v - \omega \geq v_0 - \omega = \epsilon$, so that $v_t \geq v_0 + t\epsilon$. This would contradict the fact that $(v_t, b_t) \in \bar{\mathcal{E}}$ (a compact set) for every history. Thus it must be that $v_0 \leq \omega$.

For the second point, suppose that Assumption 2 is violated. Consider any productive equilibrium σ. Dropping to an on-path history if necessary, we may assume σ doesn't start with a freeze. If $\hat{\lambda}_0 > 0$, then σ isn't an aligned equilibrium. If $\hat{\lambda}_0 < \infty$, then agent IC implies $v_0 - \tilde{v}_0 \geq \underline{\theta}$. Therefore,

$$v_0 \geq v_0 - \tilde{v}_0 \geq \underline{\theta} > \omega.$$

Appealing to the first part, it must be that σ is not an aligned equilibrium.

Under Assumption 2, $\sigma^{\bar{\tau}}$ is an equilibrium providing revenue ω, proving the third item and the other direction of the second item. □

The above result has a lot of content. The first part gives a firm upper bound on how much value can be created in an aligned equilibrium. If the principal wants the agent to behave, she has to stop him from taking bad projects. In an aligned equilibrium—in which the principal's payoffs are directly proportional to the agent's—anything that punishes the agent punishes the principal just as much. Since the rolling budget rule $\sigma^{\bar{\tau}}$ entails as little punishment as possible subject to agent IC whenever the principal is delegating, it's best for both players among any aligned equilibrium. In what follows, we refer to $\sigma^{\bar{\tau}}$ as our **Aligned Optimal Budget**.

One important consequence of the above is that aligned-equilibria cannot hope to acheive first-best for the principal, even as the players become very patient. Indeed,

$$\frac{\text{aligned optimal profit}}{\text{first-best profit}} = \frac{\omega(1 - \frac{c}{\bar{\theta}})}{\eta(\bar{\theta} - c)} = \frac{\frac{\omega}{\eta}(\bar{\theta} - c)}{(\bar{\theta} - c)\bar{\theta}} = \frac{\bar{\theta} - \underline{\theta}}{\bar{\theta}},$$

which is < 1 and doesn't vary based on how patient the players are.

5 Perennial Projects

In the current section, we showcase another instrument the principal can employ to incentivize the agent. In aligned equilibria, there being no means to reward the agent for patience, costly shutdown bore the full cost of providing incentives. There is, however, another tool at the principal's disposal: bad projects—which are readily available and provide value to the agent. The mutual expectation of a bad project t periods in the future provides a benefit of $e^{-t}\underline{\theta}$ to the agent at a cost of $e^{-t}(c - \underline{\theta})$ to the principal. As multiple bad projects can be awarded at various times, this mechanism amounts to transferable utility. Perhaps, by allowing bad projects in case the agent has enough bad luck, the principal can benefit from burning less value in case the agent has good luck. Below, we study a very simple contract that attempts to balance this tradeoff, and we compare it to the aligned optimum.

Consider the cyclic contract with fiscal year of length $T \in (0, \infty)$, in which $\mathcal{A}$ is allowed to adopt one project per fiscal year.[15] Every fiscal cycle, $\mathcal{P}$ delegates until $\mathcal{A}$ takes a project, at which point $\mathcal{P}$ freezes until the year's end. $\mathcal{A}$ best responds: he takes the first good project that comes, but resorts to a bad one if by the end of the year no good project has arrived. Rather than working with cycle length T as our parameter, it's convenient to work with the transformed variable

$$z = \sum_{k=0}^{\infty} e^{-kT} = \frac{1}{1 - e^{-T}} > 1, \text{ so that } e^{-T} = \frac{z-1}{z}.$$

We're interested in how much value the cyclic contract can deliver to the principal. Unlike with aligned equilibria, $\mathcal{P}$ and $\mathcal{A}$ may disagree on which cycle length is best. To the agent, a shorter cycle is unambiguously better. To the principal, there is a genuine tradeoff: a shorter fiscal year means more agent value, but also higher costs.

The expected discounted number of bad/good projects are

$$\begin{aligned} b(z) &= ze^{-T}e^{-\eta T} = (z-1)\left(\frac{z-1}{z}\right)^{\eta}, \text{ and} \\ g(z) &= z\int_0^T e^{-t}\eta e^{-\eta t}\, \mathrm{d}t = \frac{z\eta}{1+\eta}[1 - e^{(1+\eta)T}] \\ &= \frac{\eta}{1+\eta}(z-b). \end{aligned}$$

Which $z > 1$ maximizes $\mathcal{P}$'s profit? Letting $\rho := \frac{\bar{\theta}-c}{c-\underline{\theta}} > 0$, profit is proportional to

$$\rho g - b = \frac{\rho\eta}{1+\eta}(z-b) - b \propto \frac{\rho\eta}{\rho\eta+1+\eta}z - b.$$

It's straightforward to show that profit is uniquely[16] maximized at $z^* = z(\eta)$, the unique value $z^* > 1$ with[17]

$$b'(z^*) = \frac{\rho\eta}{\rho\eta+1+\eta}.$$

[15] i.e. One project during $((k-1)T, kT]$ for every $k \in \mathbb{N}$.

[16] Uniqueness is guaranteed by the convexity of b, together with a simple analysis of the limiting cases.

[17] The given z^* describes an equilibrium if and only if the principal has nonnegative continuation profit at the end of a project cycle. Given the below work comparing profit to first-best profit, this condition holds when η is sufficiently high.

How high is the profit, as a fraction of first-best profit? Some algebra shows that

$$\begin{aligned}\frac{\text{profit from optimal perennial project contract}}{\text{first-best profit}} &= \frac{\rho g - b}{\rho\eta - 0} \\ &= \frac{1}{1+\ell},\end{aligned}$$

where $\ell(\eta) := \frac{\eta}{z(\eta)}$. Then,

$$\left(1+\frac{-\ell}{\eta}\right)^{\eta}(1+\ell) = b'(\eta) = \frac{\rho\eta}{\rho\eta+1+\eta} \xrightarrow{\eta\longrightarrow\infty} \frac{\rho}{\rho+1} = \frac{\bar{\theta}-c}{\bar{\theta}-\underline{\theta}}.$$

Therefore, $\ell^* := \lim_{\eta\to\infty}\ell(\eta)$ exists and is the unique positive number satisfying $e^{-\ell^*}(1+\ell^*) = \dfrac{\bar{\theta}-c}{\bar{\theta}-\underline{\theta}}$. In particular $\ell^* > 0$ (and so $\frac{1}{1+\ell^*} < 1$) so that, even as the players become arbitrarily patient, the perennial project contract cannot approximate efficiency.

There are two lessons to be learned here. First, the ability to reward the agent with bad projects, following fiscal restraint, can benefit the principal. Indeed, for some parameter values the best perennial project contract outperforms[18] the aligned optimum. Second, by providing incentives one year at a time, the principal wastes her capacity to link decisions across years. Indeed, if at year's end the principal were to roll over the agent's remaining budget into the next fiscal cycle, both player would be made better off.

6 Dynamic Capital Budgeting

In the previous two sections, we presented some sensible budget rules. First, in aligned equilibria, the agent is punished via mutual money burning for each project but cannot be rewarded for fiscal restraint. Next, in perennial project contracts, in each time interval of fixed length the agent has a fixed[19] endowment of projects. Such contracts both punish the agent for taking a project and reward him for waiting. However, the reward is discontinuous, the endowment restarting at the end of each fiscal year. Such discontinuity is inefficient: conditional on reaching the budgetary cycle's end with no good project, the agent is certain to adopt a bad one. At that point, both players would love to extend the deadline: have the agent delay for an instant to keep searching for a good project, and only then (if his bad luck continues) adopt a bad project. In a sense, the principal would want the agent to save his "stock" of projects for the future. Our next contract, the **Dynamic Capital Budgeting**

[18] For any $\alpha, \beta \in (0,1)$, there are parameter values $\bar{\theta} > c > \underline{\theta} > 0$ such that $\frac{\bar{\theta}-\underline{\theta}}{\bar{\theta}} = \alpha$ and $\frac{1}{1+\ell^*} = \beta$.

[19] Up to issues of integer rounding issues.

(DCB) contract, is an attempt to implement this logic.

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 moment in time, players follow Controlled Budgeting or Capped Budgeting, depending on the agent's balance, x.

Capped Budget $(x > 0)$
The account balance grows at the interest rate $r = 1$ as long as $x < \bar{x}$. While $\mathcal{A}$'s account is in the black, $\mathcal{P}$ fully delegates project choice to $\mathcal{A}$. However every project $\mathcal{A}$ initiates reduces the account balance to $x - 1$ (whether or not the latter is positive). 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 $(x \leq 0)$
When $x < 0$, the agent is over budget, and the principal freezes for duration $\log \frac{\omega}{\omega - \underline{\theta}|x|}$. Following this freeze—more severe the further over budget is the agent—the principal delegates at balance $x = 0$. The continuation contract when the balance is $x = 0$ is a form of budget: the Aligned Optimal Budget.

The formal definition is below:

Definition. *Define the* ***Dynamic Capital Budgeting*** *contract* $\sigma^{x,\bar{x}}$ *as follows:*

1. *The* ***Capped Budget*** *regime:* $x > 0$.
   - *While* $x \in (0, \bar{x})$*:*
     - $\mathcal{P}$ *delegates, while* $\mathcal{A}$ *takes any available good projects and no bad ones;*
     - *If* $\mathcal{A}$ *picks up a project, the balance jumps from* x *to* $x - 1$*;*
     - *If* $\mathcal{A}$ *doesn't take a project,* x *drifts according to* $\dot{x} = rx = x > 0$.
   - *When* x *hits* $\bar{x}$*:*
     - $\mathcal{P}$ *delegates, and* $\mathcal{A}$ *takes a bad 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* $\log \frac{\omega}{\omega - \underline{\theta}|x|}$*;*

Then they switch to the Aligned Optimal Budget $\sigma^{\bar{\tau}}$.

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. In the Capped Budget regime, for a given account cap, the balance has non-monotonic profit implications. If the account balance runs low, there is an increased risk of imminently moving towards the low-revenue Controlled Budget regime. If the account balance runs high, the principal faces more bad projects in the near future. Note that $\mathcal{P}$'s inability to credibly commit limits how high the cap can be. If the cap were too high, the principal's continuation value would be negative at high balances, and she would rather unilaterally shut down. A Capped Budget allows our principal, who cannot credibly offer long-term compensation to her agent, to fruitfully balance punishment and reward.

The following figure represents one possible realized path of the account balance over time. Notice that bad projects are clustered: immediately after a bad project, the high balance of $\bar{x}-1$ means the next project is likely bad. Given exponential growth, this effect is stronger the higher is the cap.

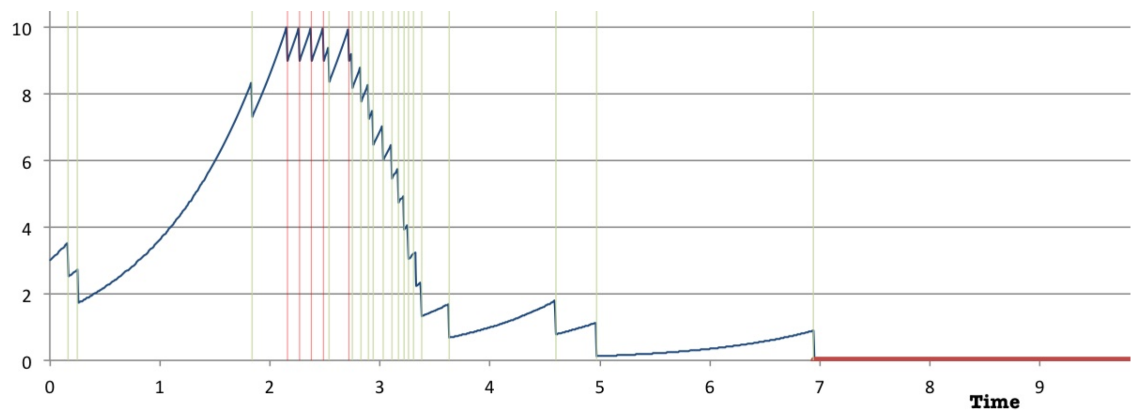

Figure 2: One realization of the balance's path under Controlled Budgeting (with $\bar{x} = 10$). Bad projects are clustered, and the account eventually runs dry.

Theorem. *Fix an account cap, $\bar{x} > 0$, and consider $\{\sigma^{x,\bar{x}}\}_{x\in[-1,\bar{x}]}$ as defined above.*

1. *Expected discounted revenue is $v(x) = \omega + \underline{\theta}x$.*

2. *Expected discounted number of bad projects is $b(x) = b^{\bar{x}}(x)$, uniquely determined by:*

   - $b|_{(-\infty,0]} = 0$
   - *b is differentiable.*

- $(1+\eta)b(x) = \eta b(x-1) + xb'(x)$ *for* $x > 0$.
- $b(\bar{x}) = 1 + b(\bar{x}-1)$.

3. *When* $x \leq 0$*:* $\sigma^{x,\bar{x}}$ *is an (aligned) equilibrium.*

4. *When* $x > 0$*:* $\sigma^{x,\bar{x}}$ *is an equilibrium if and only if*

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

Proof. The first point follows from substituting into the v promise-keeping constraint, and noting that (by work in Section 4) $\sigma^{0,\bar{x}}$ yields revenue ω.

The second point follows from our work in Section 10. The third point follows from work in Section 4.

For the fourth part, $v(x) - v(x-1) = [\omega + \underline{\theta}x] - [\omega + \underline{\theta}(x-1)] = \underline{\theta}$ 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. But now, revenue is linear, and b is (by work in Section 10) convex. Therefore profit is concave in x. So profit is nonnegative on the whole domain $[-1, \bar{x}]$ if and only if it's nonnegative at the top. □

6.1 Optimality

To gain some intuition as to why the above equilibrium should be optimal, consider how the principal might like to provide different levels of revenue. For revenue $v \leq \omega$, it's simple: we know the principal can provide said revenue efficiently, i.e. via some aligned equilibrium. We also know that other contracts—for instance, perennial project contracts—may yield higher revenue; the key issue is how to provide such higher revenue levels optimally. The principal can provide revenue and incentives via two instruments: (i) punishing the agent for spending, and (ii) rewarding the agent for fiscal restraint. Reminiscent of Ray (2002), the DCB contract backloads costly rewards as much as possible. Subject to satisfying the agent's incentive constraint, the DCB contract uses the minimal punishment possible, i.e. $v - \tilde{v} = \underline{\theta}$, whenever delegating to the agent. Increasing the punishment would accomplish two things, both of them profit-hindering:

1. Following good luck, it brings the players closer to a low-revenue continuation, where the principal will then have to inefficiently freeze.

2. In accordance with promise-keeping, the increased punishment must be paired with a reward for waiting. Following bad luck, this brings the players to a very high-revenue continuation more quickly, which will necessarily entail more bad projects.

Below is our main theorem:

Theorem 2. [20] *Let $\bar{v} \geq \omega$ denote the highest agent value attainable in equilibrium. Let $B : [0, \bar{v}] \to \mathbb{R}_+$ take any agent value v to the* efficient *number of bad projects, i.e. minimum number of (expected discounted) bad projects in any equilibrium that provides the agent said value. Then:*

$$B(v) = \begin{cases} 0 = vB'(v) & \text{if } v \in [0, \omega] \\ \dfrac{\eta B(v - \underline{\theta}) + (v - \omega)B'(v)}{1 + \eta} & \text{if } v \in (\omega, \bar{v}) \\ 1 + B(v - \underline{\theta}) & \text{if } v = \bar{v} > \omega, \end{cases}$$

for any $v \in [0, \bar{v}]$. Moreover, if $\bar{v} > \omega$, then $\pi(\bar{v}, B(\bar{v})) = 0$.

In the proof, we first analyze the discrete time game and characterize the set of equilibrium supported values $\mathcal{E}^{\Delta}$. Allowing for public randomization guarantees us the convexity of the frontier. Whenever incentivizing picky project adoption by the agent, the principal inflicts the minimum punishment possible. Next, we show that the frontier is self-generating and private mixing unnecessary. Collectively this yields a Bellman equation in discrete time. Next, we show that initial freeze is inefficient for values above ω, and that bad project adoption is wasteful, except when used to provide value $\bar{v}$. Finally we analyze the limit of the sets of equilibrium supported values $\mathcal{E}^{\Delta}$, as Δ shrinks.

Because the convex continuous function B meets the zero profit line at both endpoints, every equilibrium value profile in $\mathcal{E}^*$ is a convex combination of $\vec{0}$ and something from ($\mathcal{E}^*$'s frontier) the graph of B.

Corollary. *Suppose there exists some non-aligned equilibrium.*[21] *Then $\bar{v} > \omega$, and there is some cap $\bar{x} > 0$ such that:*

1. *For any equilibrium value profile $(v, b) \in \mathcal{E}^*$ with $v > \omega$, there is a DCB contract with some initial positive balance and cap $\bar{x}$ which provides the same revenue and bad projects $\leq b$, and therefore provides profit $\geq \pi(v, b)$.*

2. *Any equilibrium value profile $(v, b) \in \mathcal{E}^*$ can be provided exactly with temporary shutdown followed by a DCB contract with cap $\bar{x}$.*

[20] The discrete time counterpart is Theorem 6.

[21] The result isn't vaccuous. If $(1 - \frac{c}{\bar{\theta}})\omega \geq c - \underline{\theta}$, then a DCB contract with account cap and initial balance of 1 is an equilibrium.

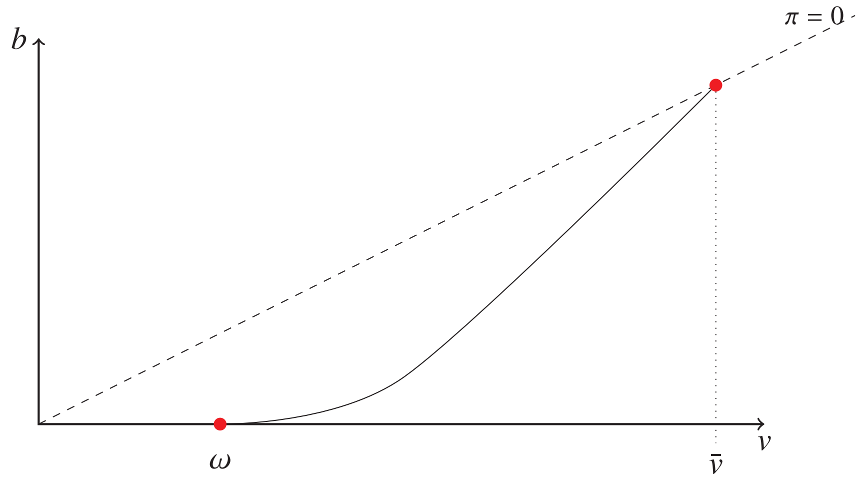

Figure 3: The solid line traces out the frontier B. The equilibrium value set is the region between B and the dashed zero profit line.

3. *There is a unique profit-maximizing profile* $(v^*, b^*) \in \mathcal{E}^*$.
   Let $x^* \in (0, \bar{x})$ *be the unique solution to*

$$\frac{d}{dx}\bigg|_{x^*} \left[\left(1 - \frac{c}{\bar{\theta}}\right)(\omega + \underline{\theta} x) - c\left(1 - \frac{\underline{\theta}}{\bar{\theta}}\right) b^{\bar{x}}(x)\right] = 0,$$

   or equivalently

$$\frac{d}{dx}\bigg|_{x^*} b^{\bar{x}}(x) = \frac{\underline{\theta}(\bar{\theta} - c)}{c(\bar{\theta} - \underline{\theta})}.$$

   Then $(v^*, b^*) = \big(\omega + \underline{\theta} x^*, b^{\bar{x}}(x^*)\big)$.

In particular, every aligned equilibrium is strictly suboptimal.

Proof. By Theorem 1, we know $\bar{v} \geq \omega$. Since there is a non-aligned equilibrium, there is some equilibrium which begins with delegation and an immediate bad project; let $\tilde{v}$ denote the continuation revenue after this bad project. If $\tilde{v} > \omega$, then obviously $\bar{v} > \omega$. If $\tilde{v} \leq \omega$, then principal participation implies $(1 - \frac{c}{\bar{\theta}})\omega \geq (c - \underline{\theta})$. In this case, the DCB contract with $\bar{x} = x = 1$ is an equilibrium, so that $\bar{v} > \omega$.

Let $\bar{x} := \frac{\bar{v} - \omega}{\underline{\theta}}$. Given Theorem 2, the first and second parts are immediate. For the third part, since the principal's objective is linear over $\mathcal{E}^*$, the optimal profit is attained on the graph of B. Since B is convex, the first-order approach suffices. By work in Section 10, B is C^1, so that the FOC holds exactly at the optimum, which can only be true for $v > \omega$. Again, by work in Section 10, B is strictly convex on $[\omega, \bar{v}]$, so that the optimum is unique. Taking an affine transformation: there is a unique optimal balance x^*, and it's > 0. □

To us, it's remarkable that the full equilibrium set can be described, even implicitly, by one scalar equation; this simple structure derives from the paucity of instruments at the principal's disposal. Even more remarkable, the DCB contract isn't just optimal; it's essentially uniquely optimal. While there is some flexibility in providing agent values below ω,[22] the optimal way to provide revenue $v \in (\omega, \bar{v})$ is unique. Given that B is strictly convex at v, $B(v) < 1 + B(v - \underline{\theta})$, and $B(v) < vB'(v)$, optimality demands initial delegation paired with picky project adoption and minimum punishment per project. Every principal-optimal contract, therefore, consists of two regimes, the first of which is Capped Budgeting (with the same cap and initial balance). In this sense, dynamic budgeting is not just a useful tool for repeated delegation but in fact a necessary one.

7 Extensions

In the current 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. It may be to the principal's advantage to provide an agent monetary reward for his fiscal constraint instead of allowing the agent to initiate a bad project. Which is more efficient depends on the relative size of the marginal cost of allowing the agent to initiate bad projects[23] $(c - \underline{\theta})B'(\bar{v})$ 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,[24] 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 Biais et al. (2010).

[22] We saw such an example in Section 4, in the "A Smaller Stick" subsection.

[23] This calculation is done using the B from our original model, as characterized in Theorem 2. The condition is correct if $\bar{v} > \omega$; otherwise, it's optimal to use monetary transfers.

[24] 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 values.[25] That is, at any moment, the principal must choose between fully delegating and ending the game forever. Naively, this constraint may seem a burden on the principal. However, given our optimal contract (with all freeze backloaded to the Controlled Budget regime), we see it changes nothing. Indeed, replacing temporary freeze with stochastic termination[26] leaves payoffs and incentives unchanged.

Agent Replacement

We consider the case in which the principal can fire the agent and immediately pick up a new agent. A fired agent gets continuation payoff of zero.[27] Every time the principal hires a new agent, she proposes a new equilibrium. In our original model, the only relevant constraint for the principal was the participation constraint. Now, the better outside option for the principal limits the credible promises she can make to the agent: she can terminate this relationship and propose a contract to a new agent. That $\mathcal{P}$'s outside option and her optimal value coincide forces the principal, unable to commit, to claim the same continuation payoff following any history. Given this, revenue can be no higher than ω.[28] Finally, notice that ω is attainable with no bad projects: $\mathcal{P}$ delegates, $\mathcal{A}$ initiates only good projects, and every project is followed with stochastic agent replacement.[29] In this contract, the principal is always delegating to the current agent, who in turn is adopting only good projects. Although each relationship has an expected revenue of ω—and thus is less profitable than in the optimal DCB contract—the principal's expected total profit across relationships is the first-best $\eta(\bar{\theta} - c)$.

[25]The agent could have a positive outside option. As long as it's below $\omega - \underline{\theta}$, the same argument holds.

[26]Keep Capped Budgeting exactly the same. In Controlled Budgeting, replace the duration $\log \frac{\omega}{\omega - \underline{\theta}|x|}$ freeze with a probability $\frac{\theta}{\omega}|x|$ termination. As the principal prefers not to terminate the relationship, the randomization must be public.

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

[28]Consider the highest agent value attainable in equilibrium. Along similar lines to the proof in Theorem 1, either the value is below ω or the agent takes an initial bad project, the latter contradicting constant principal continuation.

[29]The probability of replacement is given by $\frac{\theta}{\omega}$.

Diverse Projects

What if there are different qualities of projects? Consider the alternative model wherein project quality θ is drawn[30] from $[\underline{\theta}, \bar{\theta}]$ instead of $\{\underline{\theta}, \bar{\theta}\}$. Maintaining the assumptions that the project cost c is θ-invariant, that good projects are infinitely scarcer than bad ones, and a suitable adaptation of Assumption 2, the optimal contract will be a qualitatively similar DCB. The Controlled Budget will demand picky project adoption (now with any project of type $\theta \geq c$) and will punish any project adoption with a freeze of length τ to make the agent indifferent when the principal is. The Capped Budget contract will have the same shape as before, with one extra feature: when the balance is at the cap, projects are at a discount, costing only $\frac{\theta}{c}$ from the account.

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) in case he exerts fiscal restraint for a long enough time. With full commitment power, slight modifications of our argument show that $\bar{v}$ is the agent's first-best.[31] This case is discussed more directly in Guo and Hörner (2014) using the methodology of Spear and Srivastava (1987).

8 Final Remarks

In this paper, we have presented an infinitely repeated instance of the delegation problem. The agent won't represent the principal's interests without being offered dynamic incentives, while the principal cannot credibly commit to long-term rewards. The latter is a key ingredient of our leading applications: a university can't credibly allow its physics department one hire per day forevermore, no matter how little hiring the department has done to date.

First, we characterize equilibria that eschew reliance on lenience-based rewards. The principal's hands are tied: she can only punish the agent by limiting her future reliance on his private information, harming herself. The Aligned Optimal Budget pairs discerning project adoption with the minimum length freeze to incentivize it.

Second, we explore what efficiency gains are possible if bad projects are used as a costly bonus. The promise of future rewards can better incentivize good behavior from the agent, and the value of the future relationship can make such rewards credible for the principal. We characterize the principal-optimal such equilibrium, the Dynamic Capital Budget contract,

[30] For ease of exposition, we describe the case of full-support draws with no atom at c.

[31] We focus on the discrete time setting, so that the agent's first-best outcome is finite.

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, every project being financed from the account. The agent takes every available good project; only when at the cap, he takes a bad one. Eventually, the account runs dry; the players transition to the second regime, Controlled Budgeting, wherein they follow the Aligned Optimal Budget. Beyond maximizing the principal's value, the DCB contract traces out (as the balance varies) the whole equilibrium value set; we note that the analysis and result proceed at any fixed discount rate.[32]

The optimal contract suggests rich dynamics for the relationship. Early on, in Capped Budgeting, the relationship is highly productive but low-yield: every good project is adopted, but some bad projects are as well. The lack of principal commitment limits the magnitude of credible promises, resulting in a transient Capped Budgeting phase. As the relationship matures to Controlled Budgeting, it's high-yield but less productive: only good projects are adopted, but some good opportunities go unrealized. In this sense, the relationship drifts toward conservatism.

While our main applications concerned organizational economics outside of the firm, we believe our results also speak to this canonical setting.[33] 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.

---

[32]In particular, nothing we do is about the patient limit.

[33]The conflict of interests in our model may reflect an empire-building motive on the part of a department. Alternatively, it can be viewed as an expression of the Baumol (1968) sales-maximization principle.

9 APPENDIX: Characterizing the Equilibrium Value Set

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

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

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 Abreu, Pearce, and Stacchetti (1990), APS. To provide the players a given value $y = (v, b)$ from today onward, we factorize 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 (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 1. *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:*[34]

  1. *(Promise keeping):*

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

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

2. *(Incentive-compatibility):*

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

- *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 a.s.-$\mu(\hat{y})$.*

- *Let $W(Y) := \{y \in \mathbb{R}^2 : \ y$ is enforceable w.r.t. $Y\}$.*

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

Adapting methods from Abreu, Pearce, and Stacchetti (1990), one can readily characterize $\mathcal{E}$ via self-generation, through the following theorem.

Theorem 3. *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[35] $\mathcal{E}^* \subseteq W(Y_0) \subseteq Y_0$. 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 theorem, a complete enough understanding of 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 $\pi|_Y \geq 0\}$.*

[35]This can be ensured, for instance, by letting $\mathcal{E}^*$ contain the feasible set, scaled by $\frac{1}{1-\delta}$.

We need only work with potential value sets in $\mathcal{Y}$. Indeed, the feasible set $\bar{\mathcal{E}}$ belongs to $\mathcal{Y}$, and it's 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 above theorem that $\mathcal{E} \in \mathcal{Y}$.

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

Definition 2. *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):*

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

  2. *(Agent incentive-compatibility):*

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

  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 $\hat{W}(Y)$.*[36]

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's 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.

[36]Indeed, it's the union of δY and a projection of the compact set $\{(a, y) \in [0, 1] \times \mathbb{R}^2 : y \text{ is } a\text{-Pareto enforceable w.r.t. } Y\}$.

Lemma 2. *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's a subset of the latter by construction, and of the former by choosing $\tilde{y} = y' = \vec{0}, p = 1, \bar{a} = a = 1$, and letting $\check{y}$ range over 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 $p = 0$, $\bar{a} = 1$, and $\check{y} = \vec{0} \in Y$, it's 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[37]

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

Now suppose $p\bar{a} > 0$, and define[38] $a^p := \dfrac{a}{\bar{a}}$ and

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

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 $co\hat{W}(Y)$.[39] □

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's never a need to offer excessive punishment. That is, it's without loss to (1) make the agent's IC constraint (to resist bad projects) bind if he is being discerning, and (2) not respond to the agent's choice if he's being indiscriminate.

[37]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 $y = (1-p)\delta\check{y}$.

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

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

Lemma 3. *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):*

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

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

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

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):*

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

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's 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. Now, define

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

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

In the case of $a < 1$, define

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

By construction, $\delta[v(z') - v(\tilde{z})] = \underline{\theta}$, as desired.[40]

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. □

[40] 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 $a = 0$.

In the first part of the lemma, $\delta[v(y')-v(\tilde{y})] \in [\underline{\theta}, \bar{\theta}]$ has been replaced with $\delta[v(y')-v(\tilde{y})] = \underline{\theta}$. That is, it's 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, familiar from the Aligned Optimal Budget: without loss of generality, the principal uses the minimal possible punishment. The lemma also yields the following:

Lemma 4. *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 3, there exist $\tilde{z}, z' \in Y$ such that:

1. (Promise keeping):

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

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

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

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 $z^* \in Y$.

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

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

while

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

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

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

$\square$

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

9.3 Self-Generation for Frontiers

Through Lemmata 2, 3, and 4, 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's immediate that for $Y \in \mathcal{Y}$, the function B_Y belongs to $\mathcal{B}$.

Notation. *Define the following functions:*

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

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,[41] rather than on value sets themselves.

As it turns out, we really only need Lemmata 2, 3, and 4 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 5. *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*[42] *of* $\inf_{i\in\mathbb{I}} B_{Y_i}$.

The following theorem 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.

Theorem 4. *Fix any* $B \in \mathcal{B}$ *and* $v \in \mathbb{R}$. *Then:*

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

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

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 5. Then, appealing to Lemma 4, $T_a B \geq T_0 B$ 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 3 delivers $\tilde{z}, z' \in Y$ such that

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

[41] Since a set contains more information than its efficient frontier, passing to $\mathcal{B}$ and T may lose some information relative to $\mathcal{Y}$ and W. Our subsequent analysis shows that this is not a problem in our setting.

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

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

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

Solving for the agent values yields

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

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

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

- Consider any $y \in W_1(Y)$:
  Lemma 3 now delivers $z' = (v', b') \in Y$ such that

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

which can be rearranged to

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

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

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

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

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

□

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 Theorem 4 delivers our Bellman equation, Corollary 1. Then, in Proposition 2, we characterize the discrete time equivalent of the Aligned Optimal Budget. Finally, in Theorem 6, we fully characterize the frontier $B_{\mathcal{E}^*}$.

Theorem 5. *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 2, $y \in W_f(Y) \cup \bigcup_{a\in[0,1]} W_a(Y)$.

- 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

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

  Then

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

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

Next, consider any $v \in \text{dom}(B_Y)$. There is some probability measure μ on the extreme points of Y such that $(v, B_Y(v)) = \int_Y y \, \mathrm{d}\mu(y)$. 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, $(v, B_Y(v)) = \int_{v(Y)} (u, B_Y(u)) \, \mathrm{d}\mu_1(u)$, so that

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

where the last inequality follows from Jensen's theorem.

This completes the proof.

□

The Bellman equation follows immediately.

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

In line with the continuous time model, we characterize the payoffs attainable in equilibria with no bad projects.

Notation. *Define the discrete time[43] marginal value of restraint, $\omega := \dfrac{h}{1-\delta}(\bar{\theta} - \underline{\theta})$.*

Proposition 2 (Aligned Optimum)**.**
Suppose $\delta\omega \ge \underline{\theta}$.[44] Given $\hat{v} \in \mathbb{R}_+$, the following are equivalent:

1. *There is some self-generating $B \in \mathcal{B}$ with $B(\hat{v}) = 0$.*

2. *The efficient frontier $B_{\mathcal{E}^*}$ of the PPE set satisfies $B_{\mathcal{E}^*}(\hat{v}) = 0$.*

[43] Notice that, given parametrization $(\delta, h) = (1 - \Delta, \eta\Delta)$, this coincides exactly with the definition from the main text.

[44] This is a discrete time expression of Assumption 2.

3. $\hat{v} \le \omega$.

Proof. That (1) implies (2) follows from self-generation above.

Now 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

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

Finally, if $TB_{\mathcal{E}^*}(v) = T_0 B_{\mathcal{E}^*}(v)$, then[45]

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

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

Now if (2) holds and (3) doesn't, applying the above inductively yields a sequence[46] $v_n \to \infty$ on which $B_{\mathcal{E}^*}$ takes value zero. This would contradict the compactness of $B_{\mathcal{E}^*}$'s domain.

Finally, suppose (3) holds, and consider $B \in \mathcal{B}$ given by $B(v) = \begin{cases} 0 & \text{if } v \in [0,\omega] \\ \infty & \text{otherwise.} \end{cases}$

$B(\hat{v}) = 0$ by (3), and B is self-generating since $B = T_0 B$. This verifies (1). □

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 6. *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 $\underline{v} \ge \omega$ such that*

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

Proof. $B(\omega) = 0$, and B is convex. Therefore, B is strictly increasing above ω on its domain, so that[47] $\check{T}_f B(v) > \check{T}_0 B(v)$, confirming the first point.

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

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

[47] The relationship is as shown if $\check{T}_0 B(v) < \infty$. Otherwise, $T_f B(v) = T_0 B(v) = \infty$.

Given v,

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

is a nonnegative constant.

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

$$v\mapsto \check{T}_0B(v)-\check{T}_1B(v)=\delta(1-h)\left[B\left(\frac{v-[\theta_E-\underline{\theta}]}{\delta}\right)-B\left(\frac{v-\theta_E}{\delta}\right)\right]-(1-h)$$

is increasing on its proper domain. The second point follows.[48] □

Theorem 6 (Equilibrium Frontier)**.**
Suppose $\delta\omega\geq\underline{\theta}$. Let $B:=B_{\mathcal{E}^*}$ and $\bar{v}:=\max dom(B)$.

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

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

$$\begin{aligned} B(v) &= T_0B(v) \textit{ for } v\in\left[\omega,\ \ \delta\bar{v}+(\theta_E-\underline{\theta})\right];\\ B(v) &\textit{ is } \textit{ affine in } v \textit{ for } v\in\left[\delta\bar{v}+(\theta_E-\underline{\theta}),\ \ \bar{v}\right];\\ B(\bar{v}) &= T_1B(\bar{v}). \end{aligned}$$

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

Proof. The first point follows directly from Lemma 2. Now suppose $\bar{v}>\omega$.
Let $\underline{v}:=\delta\bar{v}+(\theta_E-\underline{\theta})$. Any $v>\underline{v}$ has $\frac{v-[\theta_E-\underline{\theta}]}{\delta}>\frac{\underline{v}-[\theta_E-\underline{\theta}]}{\delta}=\bar{v}$, so that $B\left(\frac{v-[\theta_E-\underline{\theta}]}{\delta}\right)=\infty$, and therefore (appealing to Theorem 4) $T_0B(v)=\infty$. Therefore, the cutoff defined in Lemma 6 is $\leq\bar{v}$.

Since T_0B, T_1B are both convex, there exist some $v_0,v_1\in[\omega,\bar{v}]$ such that $v_0\leq v_1,\underline{v}$, and:

$$\begin{aligned} B(v) &= 0 \text{ for } v\in[0,\omega];\\ B(v) &= T_0B(v) \text{ for } v\in[\omega,v_0]\,;\\ B(v) &\text{ is } \text{ affine in } v \text{ for } v\in[v_0,v_1]\,;\\ B(v) &= T_1B(v) \text{ for } v\in[v_1,\bar{v}]\,. \end{aligned}$$

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. $\supseteq$) such that the above decomposition is still correct.

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

- If $\frac{v_1-\theta_E}{\delta}\geq v_0$, then (appealing to Theorem 4) the right-side derivative $T_1B'=m$ in some neighborhood of v_1 in $[v_1,\bar{v}]$. By convexity of B, this would imply $B(v)=T_1B(v)=B(v_1)+m(v-v_1)$ in said neighborhood. Then, by maximality of v_1, it must be that $v_1=\bar{v}$.

[48] Because wherever $\check{T}_iB(v)\geq\check{T}_jB(v)$, we have $T_iB(v)\geq T_jB(v)$ as well.

- If $\frac{v_1-\theta_E}{\delta} < v_0$, then minimality of v_0 implies (again using Theorem 4) 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 $\frac{v_0-[\theta_E-\underline{\theta}]}{\delta} < \underline{v}$), then Theorem 4 gives us

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

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. □

9.5 The Efficient Frontier: Continuous Time

In this section we proceed to take the limit of the discrete time equilibrium value set frontier. We consider the limit with $(\delta, h) = (\delta_\Delta, h_\Delta) := (1-\Delta, \eta\Delta)$ (and so $\theta_E = \underline{\theta} + \omega\Delta$) as the period length $\Delta \to 0$.

Define D_0^Δ and F_0^Δ on $\mathcal{B}$ via:

$$
\begin{aligned}
D_0^\Delta B &:= \frac{\eta}{h}\left[B - T_{\Delta,f}^{-1} T_{\Delta,0} B\right], \text{ and} \\
F_0^\Delta B &:= \frac{\eta}{h}\left[B - T_{\Delta,f}^{-1} B\right].
\end{aligned}
$$

Notice that

$$
\begin{aligned}
\frac{h}{\eta} F_0^\Delta B(v) &= B(v) - hB(v-\theta_E) - (1-h)B\left(v - [\theta_E - \underline{\theta}]\right) \\
&= h\left[B(v) - B(v-\theta_E)\right] + (1-h)\left[B(v) - B\left(v - \frac{h}{1-h}[\bar{\theta} - \theta_E]\right)\right] \\
\implies F_0^\Delta B(v) &= \eta\left[B(v) - B(v-\theta_E)\right] + \frac{\eta(1-h)}{h}\left[B(v) - B\left(v - \frac{h}{\eta(1-h)}\omega\right)\right],
\end{aligned}
$$

which increases to $\eta\left[B(v) - B(v-\theta_E)\right] + \omega B'(v)$ as $\Delta \to 0$.

$$
\begin{aligned}
D_0 B(v) &= \frac{B(v) - T^{-1}_{\Delta,f} B(v)}{\Delta} \\
&= \frac{B(v) - \frac{1}{1-\Delta} B\left((1-\Delta)v\right)}{\Delta} \\
&= \frac{1}{1-\Delta}\left[\frac{B(v) - B(v-\Delta v)}{\Delta} - B(v)\right],
\end{aligned}
$$

which converges to $vB'(v) - B(v)$ as $\Delta \to 0$.

If $T_0 B(v) = B(v)$, then $D^\Delta_0 B(v) = F^\Delta_0 B(v)$, so that

$$\eta\left[B(v) - B(v-\theta_E)\right] + \omega B'(v) = vB'(v) - B(v).$$

Let $B_\Delta \in \mathcal{B}$ be the frontier of the discrete time equilibrium set $\mathcal{E}^\Delta$, and let $\bar{v}_\Delta$ be the highest agent value. Let $\bar{v} = \lim_{\Delta\to 0} \bar{v}_\Delta$.

Take any $v \in [0, \omega]$. $B_\Delta(v) \to 0$ as $\Delta \to 0$ by Proposition 2 and because $\omega \to \omega$.

Take any $v \in (\omega, \bar{v})$. For sufficiently small Δ, $v \in (\omega_{\Delta,d}, \delta_\Delta \bar{v}_\Delta + (\theta_E - \underline{\theta}_\Delta))$, so that Theorem 6 tells us $T_{\Delta,0} B(v) = B(v)$. Adapting the above argument therefore tells us

$$\eta\left[B_\Delta(v) - B_\Delta(v-\theta_E)\right] + \omega B'_\Delta(v) - \left[vB'_\Delta(v) - B_\Delta(v)\right]$$

converges to zero.

Finally, appealing again to Theorem 6, $B_\Delta(\bar{v}_\Delta) - B_\Delta(\bar{v}_\Delta - \theta_E)$ converges to 1.

Applying Theorem 6 yields the following continuous time limit for $B = B_{\mathcal{E}}$:

$$
\begin{aligned}
B(v) &= \begin{cases} 0 & \text{if } v \in [0,\omega] \\ \dfrac{\eta B(v-\underline{\theta}) + (v-\omega)B'(v)}{1+\eta} & \text{if } v \in (\omega, \bar{v}) \\ 1 + B(v-\underline{\theta}) & \text{if } v = \bar{v} > \omega, \end{cases} \\
\omega &= \eta(\bar{\theta} - \underline{\theta}) \\
\bar{v} &= \max \operatorname{dom}(B) \\
\pi\left(\bar{v}, B(\bar{v})\right) &= 0.
\end{aligned}
$$

Summarizing, there is some $\bar{v} \geq \omega$ such that:

- The highest revenue attainable in equilibrium is $\bar{v}$.
- An optimal way to provide revenues $v \in [0, \omega]$ in equilibrium is with shutdown of duration $\log \frac{\omega}{v}$ followed by the Aligned Optimal Budget.

- If $\bar{v} > \omega$, then
  - The optimal way to provide revenue $v \in (\omega, \bar{v})$ in equilibrium is delegation with picky project adoption, jumping to continuation revenue $v - \underline{\theta}$ following a project, and revenue drifting according to $\dot{v} = v - \omega$ conditional on no projects.
  - The only way to provide revenue $\bar{v}$ in equilibrium is with an immediate bad project followed by providing revenue $\bar{v} - \underline{\theta}$.
  - The optimal equilibrium profit from providing $\bar{v}$ is zero.

10 APPENDIX: Delayed Differential Equation

Taking a change of variables, from agent value v to account balance $x = \frac{v-\omega}{\theta_E}$, 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 7. *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 $b = b^{(1)}$, then

1. *$b(x) = x^{1+\eta}$ for $x \in [0, 1]$.*
2. *b is C^1 on $(-\infty, 1)$ and twice-differentiable on $(0, \infty)$.*
3. *b is strictly convex on $(0, \infty)$ and convex on $\mathbb{R}$. In particular, b is unbounded.*
4. *b is strictly increasing and strictly log-concave on $(0, \infty)$. In particular, b doesn't explode in finite time.*

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've 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 $b''(\bar{x}) = \frac{\eta}{\bar{x}}[b'(x) - b'(x-1)] > 0$, 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's 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. ◻

Proposition 3. *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 decreasing 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 $\frac{b(\bar{x})}{b(\bar{x}-1)}$ 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 $\bar{x}$. □

References

Abreu, D., P. Milgrom, and D. Pearce. 1991. "Information and Timing in Repeated Partnerships." *Econometrica* 59 (6):1713–33.

Abreu, D., D. Pearce, and E. Stacchetti. 1990. "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring." *Econometrica* 58 (5):1041–63.

Alonso, Ricardo and Niko Matouschek. 2007. "Relational delegation." *The RAND Journal of Economics* 38 (4):1070–1089.

Ambrus, A. and G. Egorov. 2013. "Delegation and Nonmonetary Incentives." Working paper, Duke University & Northwestern Kellogg.

Armstrong, M. and J. Vickers. 2010. "A Model of Delegated Project Choice." *Econometrica* 78 (1):213–44.

Baumol, W. 1968. "On the Theory of Oligopoly." *Economica* 25 (99):187–98.

Biais, B., T. Mariotti, J. Rochet, and S. Villeneuve. 2010. "Large Risks, Limited Liability, and Dynamic Moral Hazard." *Econometrica* 78 (1):73–118.

Casella, A. 2005. "Storable Votes." *Games and Economic Behavior* 51:391–419.

Chassang, S. 2010. "Building Routines: Learning, Cooperation, and the Dynamics of Incomplete Relational Contracts." *American Economic Review* 100 (1):448–65.

Clementi, G. and H. Hopenhayn. 2006. "A Theory of Financing Constraints and Firm Dynamics." *The Quarterly Journal of Economics* 121 (1):1131–1175.

Espino, Emilio, Julian Kozlowski, and Juan M Sanchez. 2013. "Too big to cheat: efficiency and investment in partnerships." *Federal Reserve Bank of St. Louis Working Paper Series* (2013-001).

Frankel, A. 2011. "Discounted Quotas." Working paper, Stanford University.

———. 2013. "Delegating Multiple Decisions." Working paper, Chicago Booth.

———. 2014. "Aligned Delegation." *American Economic Review* 104 (1):66–83.

Guo, Y. 2014. "Dynamic Delegation of Experimentation." Working paper, Northwestern University.

Guo, Y. and J. Hörner. 2014. "Dynamic Mechanisms without Money." Working paper, Northwestern University & Yale University.

Hauser, C. and H. Hopenhayn. 2008. “Trading Favors: Optimal Exchange and Forgiveness.” Working paper, Collegio Carlo Alberto & UCLA.

Holmström, B. 1984. “On the Theory of Delegation.” In *Bayesian Models in Economic Theory*. New York: North-Holland, 115–41.

Jackson, M. and H. Sonnenschein. 2007. “Overcoming Incentive Constraints by Linking Decisions.” *Econometrica* 75 (1):241–57.

Levin, Jonathan. 2003. “Relational incentive contracts.” *The American Economic Review* 93 (3):835–857.

Li, F., N. Matoushek, and M. Powell. 2014. “The Burden of Past Promises.” Working paper, Northwestern University.

Li, Jin and Niko Matouschek. 2013. “Managing conflicts in relational contracts.” *the American Economic Review* 103 (6):2328–2351.

Mailath, G. and L. Samuelson. 2006. *Repeated Games and Reputations: Long-Run Relationships*. New-York: Oxford University Press.

Malcomson, James M. 2010. “Relational Incentive Contracts.” .

Malenko, A. 2013. “Optimal Dynamic Capital Budgeting.” Working paper, MIT.

Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green. 1995. *Microeconomic theory*. Oxford: Oxford University Press.

Möbius, M. 2001. “Trading Favors.” Working paper, Harvard University.

Pearce, David G and Ennio Stacchetti. 1998. “The interaction of implicit and explicit contracts in repeated agency.” *Games and Economic Behavior* 23 (1):75–96.

Ray, Debraj. 2002. “The Time Structure of Self-Enforcing Agreements.” *Econometrica* 70 (2):547–582.

Spear, S. and S. Srivastava. 1987. “On Repeated Moral Hazard with Discounting.” *Review of Economic Studies* 54 (4):599–617.