

Fostering collaboration

JOYEE DEB

Department of Economics, New York University

ADITYA V. KUVALEAR

Department of Economics, University of Essex

ELLIOT LIPNOWSKI

Department of Economics, Yale University

We study project selection and development by a principal, interacting with two agents, each of whom wants his respective project selected. When the best choice is uncertain, keeping both projects alive gives the principal the ability to adapt her choice in the future, but implies an efficiency loss of effort being spent on the project finally not chosen. We show that a time-varying threshold rule is uniquely optimal: the principal selects the first project to achieve a sufficient lead. The optimum entails initial competition, always followed by permanent collaboration. Our proof uses martingale time-change methods that apply weak solutions.

KEYWORDS. Project selection, internal competition, team production, collaboration, mechanism design without transfers.

JEL CLASSIFICATION. D72, D73, D82, D86.

1. INTRODUCTION

When faced with a choice between different courses of action, organizations often create internal competition: allowing multiple teams to develop competing approaches to solve the same problem. Keeping competing approaches alive provides option value

Joyee Deb: joyee.deb@nyu.edu

Aditya V. Kuvalekar: a.kuvalekar@essex.ac.uk

Elliot Lipnowski: elliott.lipnowski@yale.edu

We would like to thank Avi Acharya, Heski Bar-Isaac, Simon Board, Ale Bonatti, Francesc Dilmé, Bob Gibbons, Marina Halac, Navin Kartik, Aaron Kolb, Annie Liang, MathOverflow user #143907, Stephen Morris, Barry Nalebuff, Harry Pei, Ben Polak, João Ramos, Debraj Ray, Tomasz Szadzik, Rani Spiegler, Balázs Szentes, Yu Fu Wong, and Bill Zame for helpful feedback. We are especially grateful to Mykhaylo Shkolnikov for his advice. We also thank various conference and seminar participants, especially Martin Cripps and Meg Meyer, who served as discussants for this paper in the ERC 2019 Frontiers in Design Conference and the NBER Fall 2020 Organizational Economics Working Group, respectively. A previous draft of this paper was circulated with the title “Curtailling Competition.” Aditya Kuvalekar gratefully acknowledges financial support by MICIN/AEI/10.13039/501100011033 (Grant PGC2018-096159-B-I00) from the National Plan, Grant PID2019-107259RB-I00 from Agencia Estatal de Investigación del Gobierno de España, and MICIN/AEI/10.13039/501100011033 (Grant PID2021-123747NA-I00), which he received while he was at Universidad Carlos III de Madrid.

when the best course of action is initially uncertain. Consider the example of the information technology infrastructure firm Telstar Communications that had two distinct 50-person teams working on two competing middleware technology platforms: AX and EX (see Birkinshaw (2001)). Each team worked on its own platform, knowing the firm would ultimately adopt exactly one.¹ When the most promising approach is not clear ex ante, allowing teams to continue working on competing approaches rather than collaborating on a single approach gives an organization the ability to adapt its choice based on what the future will look like. However, this adaptive benefit must be balanced with the efficiency loss of wasted productive effort on the “wrong” approach. This tension is at the heart of our paper. We characterize an optimal selection rule that harnesses optimally the adaptive benefits of continuing to develop competing approaches and the efficiency benefits of collaborating on a single approach.

Formally, we study a finite-horizon game in continuous time, in which a principal interacts with two agents until a deadline T . Each agent has his own project. The principal evaluates the projects as they are developed and must pick one of them when the deadline arrives. At every instant, an agent decides how to allocate a unit of effort between working on his own project and providing assistance (collaborating) on the other agent's project. Effort is costless. Each project's value evolves as a Brownian motion with a drift. The drift is increasing in the total effort expended on the project by the two agents, while the Brownian shocks are exogenous. The vector of the projects' current state of development is publicly observed by both agents and the principal in real time, even though effort choices are not observed—or at least not contractible. The principal's payoff is equal to the state of the project she chooses at the deadline; she does not benefit from the state of the other project. The agents have conflicting interests, in that each wants his own project to be chosen. Since effort is costless, an agent may be willing to collaborate (work on the other agent's project) if doing so does not undermine his own chance of being selected. Our goal is to characterize the principal-optimal selection rule (that chooses a project at the deadline as a function of the history of the projects' evolution), assuming the principal can commit to any history-dependent rule.

We first note that in the principal's first-best policy, ignoring agency problems, the principal would simply wait until the deadline to see the realization of all shocks, and then pick the project with the higher final state. Moreover, at any instant, she would like both agents to collaborate on the project that currently has the higher state. This effort allocation maximizes the likelihood that effort is productively useful. The first-best policy captures the intuition that the principal wants to foster collaboration (have both agents work on the same project), while constantly adjusting the project choice to their shocks, to ensure the agents collaborate on the “right” project. However, it is not incentive compatible for the agents, because they will each strictly prefer to work on their own project. Indeed, consider another benchmark, in which the principal can make no prior commitments regarding the eventual project choice. It is easy to see that, in this case too, the principal will pick the project with the higher ex post state at the deadline.

¹IBM similarly fosters competition between teams for product ideas, encouraging different teams to try competing approaches to the same problem (Peters and Waterman (2003)).

A unique equilibrium between the agents ensues, with no collaboration. Agents would rather compete, each focusing all effort on his own project. The result is an inefficient use of the agents' efforts.

The natural question, then, is whether the principal can curtail competition and foster some efficient collaboration. Giving the principal commitment power can help. For example, suppose the principal commits to choosing the project that is first to take a lead by a specified margin. It is easy to see how such a rule can outperform the outcome under no commitment (pure competition). Moreover, it is incentive compatible, because once an agent's project is chosen irreversibly, agents are indifferent and, thus, willing to collaborate. At this point, one might wonder whether the only way for the principal to induce agents to collaborate is to make an irreversible choice at some point. One may reasonably conjecture that with strictly opposing interests, if an agent thought his project could be adopted in the future after enough improvement, however small the likelihood, he would strictly prefer to work on his own project. Perhaps surprisingly, this conjecture is false. Many incentive-compatible decision rules induce agents to switch back and forth between competing and collaborating.² The space of all incentive-compatible selection rules is rich, and our main result characterizes the optimal such rule.

We show that the unique principal-optimal selection rule has a simple form: the principal commits to a time-dependent, decreasing lead threshold $\{\hat{z}_t\}_{t \in [0, T]}$ that decreases to 0 as the deadline approaches, such that a project is chosen at the first instant t at which its state exceeds that of the other by at least $\hat{z}_t$. Agent equilibrium behavior therefore also has a simple pattern, namely, that in an initial competitive regime (before a project is chosen by the principal), agents allocate all effort toward their own project. This phase is followed by permanent collaboration; that is, agents collaborate on the chosen project (the first to achieve the threshold lead) until the deadline. In particular, regardless of the time horizon, a nondegenerate phase of collaboration always exists. The main force behind this two-phase structure is the option value arising from front-loading competition: sustaining temporary collaboration before additional competition necessitates that the principal not adapt her project choice to projects' shocks during the collaboration phase. Leaving collaboration for later (after having determined which of the two projects is likely to be chosen) reduces the likelihood of mistakes; i.e., collaborating on a project that (due to worse exogenous shocks) is not ultimately chosen. The decreasing lead threshold captures the diminishing benefits of option value from competition.

This two-phase optimal contract with inevitable collaboration from some point onward is broadly consistent with what we observe in our motivating organizational setting. For example, at Telstar, top-level executives finally chose EX over AX, and both teams subsequently collaborated on EX to build a common platform.³

The derivation of the optimal contract proceeds in four logical steps. We first show that any contract is outperformed by one in which the principal resolves her decision

²See Section 3.3 for a more detailed discussion on this point.

³Similarly, at IBM, teams are allowed to work on disparate approaches until, at some point, the firm conducts performance "shootouts" to pick one (Peters and Waterman (2003)).

quickly. Intuitively, if she does not adapt her current choice adequately to project shocks for some time, she is not utilizing the option value of competition. Speeding up her decision makes room for additional collaboration targeted to the eventual chosen project.⁴ Second, we show the principal conditions only on relative performance, because aggregate shocks are both irrelevant to the principal's objective and uninformative about agent behavior. Third, we show it is optimal for the principal to consider two-phase policies with initial competition until a stopping time, followed by a permanent switch to collaboration; in particular, she makes a constrained-efficient choice with the partial information available when ending competition, and has both agents collaborate thereafter on the chosen winner. Finally, we show that the principal optimally chooses a project only when its lead over the other project is sufficiently large, lowering her standards closer to the deadline. Intuitively, when ending competition and choosing a project, the principal foregoes the option value from adjusting to projects' future shocks, but this option value vanishes as the deadline approaches.

A notable feature of the optimal policy is that competition is always temporary. For any horizon, collaboration starts strictly before the deadline with probability 1. This implies arbitrarily large ex post inefficiencies can occur on path.

Finally, one of our main contributions is a methodological one. Given our finite horizon, controlled volatility (with zero volatility being feasible), and discontinuous flow payoffs, we cannot adopt the standard approach (à la [Sannikov \(2008\)](#)) of heuristically deriving the Hamilton–Jacobi–Bellman (HJB) equation, establishing the existence of a smooth solution, and appealing to a verification theorem. As we demonstrate, using time-change methods to analyze weak controls enables tractable value calculations for intuitive intertemporal policy modifications. We are hopeful that this approach will be useful more broadly in economic theory.

1.1 *Related literature*

At a high level, our paper begins with the premise of [March \(1962\)](#) and [Cyert and March \(1963\)](#) that individuals within organizations often have goals that are distinct from those of the organization, and the executive acts like a political broker who cannot solve such problems by simple payments.⁵

Our paper is related to work on multi-agent experimentation, e.g., [Bolton and Harris \(1999\)](#), [Keller, Rady, and Cripps \(2005\)](#), [Bonatti and Hörner \(2011\)](#), and [Halac, Kartik, and Liu \(2017\)](#). However, the trade-offs are fundamentally different. Those papers share two key features: Agents want to free-ride on each others' costly experimentation,⁶ and

⁴A similar force manifests in delayed investment when firms face uncertainty about an impending government policy choice (see [Stokey \(2016\)](#)). In that setting, as in ours, the flow of decision-relevant information is exogenous to current investment decisions.

⁵See [Gibbons \(2020\)](#) for a detailed survey.

⁶Free-riding in teams is an extensively studied topic outside the experimentation framework as well. For example, see [Holmström \(1982\)](#), [Mookherjee \(1984\)](#), and [Legros and Matsushima \(1991\)](#). More thematically related to our work, [Marino and Zabojnik \(2004\)](#) show how internal competition can be beneficial in addressing the free-rider problem.

each agent trades off the information from exploration against the myopic value of exploitation. Taken literally, our paper has neither of these features. Free-riding is completely absent, and information arrives *exogenously*. Nevertheless, because incentivizing agents to collaborate on the principal's preferred project requires contemporaneous information to be ignored, she still faces a trade-off between option value and myopic optimization—just like the exploration–exploitation trade-off in experimentation models. A particularly related experimentation model is [Durandard \(2023\)](#), which studies the experimentation problem of a principal choosing between agents. Like our model, his does not feature free-riding. [Durandard's \(2023\)](#) principal faces a bandit problem with strategic arms, trading off the option value of exploration against motivating effort. In our setting, the principal does not need to motivate (total) effort and again does not control the learning process. Rather, when the principal wants to use the exogenously generated information to inform future choices, she distorts the agents' current effort choice. Despite the significant modeling differences between our setting and those of these experimentation papers, our analysis shows our principal's problem reduces to a certain single-agent experimentation problem.⁷

The closest work to ours is that of [Bonatti and Rantakari \(2016\)](#), where each agent first chooses what type of project to develop and how hard to work on it over time, after which he negotiates over the adoption choice. Agent interests are partially aligned. A key lesson is that the project selection mechanism can feed into the development stage when agents may distort the organization's decision.⁸ This lesson sets the stage for our design problem.

Finally, we contribute methodologically to the literature on dynamic mechanism design without transfers (e.g., [Aghion and Jackson \(2016\)](#), [Deb, Pai, and Said \(2018\)](#), [Guo and Hörner \(2020\)](#)). Like us, [McClellan \(2022\)](#) employs tools from dynamic contracting in continuous time (e.g., [Sannikov \(2008\)](#), [DeMarzo and Sannikov \(2006\)](#)) to study delegated experimentation. We hope our techniques—appealing to martingale methods rather than HJB equations to simplify a volatility control problem, and passing between weak and strong solutions—can be used in future work.

2. MODEL

A principal interacts with two agents $i \in I = \{-1, 1\}$ in continuous time over a finite horizon of length T . Each agent i has a project with evolving state X^i . The principal must pick one of the two projects at the deadline T . At every instant, each agent must allocate a unit of effort between working on his own project and providing assistance on the other agent's project. Effort is costless and contributes to projects' development continuously over time. Let $a_t^i \in [0, 1]$ denote the effort that agent i allocates to his own project

⁷Less related is the literature on dynamic contests, which focuses on moral hazard and completely abstracts away from collaboration (e.g., [Benkert and Letina \(2020\)](#), [Ryvkin \(2022\)](#), [Moscarini and Smith \(2007\)](#)). In spite of this key difference in underlying incentives, it is an outcome of our analysis that the optimal selection rule induces an initial contest.

⁸This feature also arises in [Hirsch and Shotts \(2015\)](#) and [Callander and Harstad \(2015\)](#). See also [Farrell and Simcoe \(2012\)](#), who study related distortions in standards adoption across firms that produce complementary products.

at time t , having observed both projects' evolution to date. Agent i allocates the remaining $(1 - a_t^i)$ of his effort to helping agent $-i$ on his project. We interpret $(1 - a_t^i)$ as the extent to which agent i collaborates. Specifically, the productive state of each project X_t^i as of time t evolves via

$$dX_t^i = [\beta + \mu(a_t^i + (1 - a_t^{-i}))] dt + \sigma dB_t^i,$$

where B^1 and B^{-1} are independent standard Brownian motions on a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P})$ satisfying the usual conditions: $\{\mathcal{F}_t\}_{t \geq 0}$ is the natural filtration of $\mathcal{F}$ with respect to $(B_t^1, B_t^{-1})_{t \geq 0}$, parameters $\beta, \mu, \sigma \in \mathbb{R}$ have $\mu, \sigma > 0$, agent i chooses a progressively measurable $[0, 1]$ -valued stochastic process a^i on $(\Omega, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P})$, and project i has exogenous initial state $X_0^i \in \mathbb{R}$.⁹ The vector of project states is publicly observed by both agents and the principal, and effort-allocation choices are not observed. For much of our analysis, rather than working directly with the state of each project, we find it convenient to instead use a different basis. Define the *relative* and *aggregate* shock processes by

$$\Delta X := X^1 - X^{-1} \quad \text{and} \quad \Sigma X := X^1 + X^{-1},$$

and define $\Delta B, \Sigma B$, and Δa analogously. Thus, the law of motion of ΔX_t and ΣX_t is

$$d\Delta X_t = 2\mu\Delta a_t dt + \sigma d\Delta B_t,$$

$$d\Sigma X_t = 2(\beta + \mu) dt + \sigma d\Sigma B_t.$$

In our main analysis, the principal commits at time zero to an arbitrary project selection rule. Formally, the principal chooses a $\{-1, 1\}$ -valued random variable y on $(\Omega, \mathcal{F}, \mathbb{P})$ for a payoff of X_T^y ; that is, the principal's profit is equal to the productive state of the chosen project. Without loss, let us normalize $\sigma = \mu = 1$, $\beta = -1$, and $X_0^1 + X_0^{-1} = 0$.¹⁰ Therefore, after normalization, project X^i follows

$$dX_t^i = i\Delta a_t dt + dB_t^i.$$

Hence, the principal's expected payoff is

$$\mathbb{E}\left[\frac{1+y}{2}X_T^1 + \frac{1-y}{2}X_T^{-1}\right] = \frac{1}{2}\mathbb{E}[y\Delta X_T].$$

Each agent wants his project to be chosen; that is, agent i gets payoff iy . Given any (y, a^1, a^{-1}) , we can define $q_t^i := \mathbb{E}[iy|\mathcal{F}_t]$ as agent i 's continuation value at any $t \leq T$. In

⁹An agent i strategy is a measurable function (when the space $C[0, T]$ has its Skorokhod topology) $a^i : (C[0, T])^2 \times [0, T] \rightarrow [0, 1]$ with the property that $a_t^i((X_s^1, X_s^{-1})_{s \in [0, T]})$ is a function only of t and $(X_s^1, X_s^{-1})_{s \in [0, t]}$. A principal strategy is a measurable function $y : (C[0, T])^2 \rightarrow \{-1, 1\}$. Because $X_t^i - B_t^i = \int_0^t [\beta + \mu(a_s^i + (1 - a_s^{-i}))] ds$ is then a progressively measurable process on $(\Omega, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P})$ for each i , it follows that $\{\mathcal{F}_t\}_{t \geq 0}$ is also the natural filtration of $\mathcal{F}$ with respect to $(X_t^1, X_t^{-1})_{t \geq 0}$.

¹⁰Counting time in different units, we may assume without loss that $\sigma = \mu$. Then positively rescaling the state (thus, principal payoffs), we may take $\sigma = \mu = 1$. Because adding a constant to both projects' initial states or both projects' drifts simply adds a constant to the principal's payoff, we may further assume $\beta = -1$ and $X_0^1 + X_0^{-1} = 0$.

what follows, we write agent incentives from the point of view of agent 1. So we drop the superscript i and define $q_t := q_t^1 = \mathbb{E}[y|\mathcal{F}_t]$. We denote the current leader at any time $t \leq T$ by $\ell_t := \operatorname{argmax}_{i \in I} X_t^i$.

2.1 Discussion of assumptions

Unobservable effort choice. We assume the agent's choice of effort allocation is unobservable to the principal or at least is non-contractible. This is quite likely to be the case with creative effort. Many modern technology companies foster innovation and develop new solutions by giving employees the flexibility to spend a fraction of their time on projects of their choice. We view effort in our model as such unmonitored time that agents can choose to use on their own or another agent's pet project.

Agency. We study a principal–agent problem in which an incentive conflict causes the principal to distort her selection rule. A simple alternative model, but without conflict, that would generate similar dynamics is one in which a monolithic organization operates two projects at fixed flow costs and can permanently shut one down at any time. Such a model would generate identical dynamics: In fact, a crucial step in our analysis (Lemma 3) establishes that our principal's problem with agency conflict reduces to the problem mentioned above; i.e., choosing the time for irreversible project termination. In this sense, the characterization of optimal project selection in the simpler alternative model is a by-product of our analysis. Our richer modeling of an organization provides a microfoundation of how such dynamics can arise: The cost of maintaining option value arises from a collaborative agency cost of responding to contemporaneous shocks, and irreversible termination emerges from the optimal front-loading of competition.

Fixed agent payoff, no transfers. We assume an agent's payoff is independent of the chosen project's state. While somewhat extreme, this assumption captures agents with empire-building motives: Individuals increase their stature and influence within organizations when their project is chosen. The absence of transfers is assumed also in canonical models of delegated decision-making, and reflects the view that not all organizational conflict can be contracted away.¹¹ If we allowed the principal to offer agents a share of the chosen project or to make transfers, inducing collaboration would be easier. Our optimal rule highlights how, despite a severe lack of instruments, the principal can (and will) foster some collaboration.

Exogenous shocks. We assume that the shocks governing the projects' evolution are exogenous. This assumption is appropriate for settings in which effort is mainly instrumental. For instance, employees working on developing a new technology can affect their chances of success by devoting more or less effort. However, such an effort will not affect industry-level shocks that might affect the feasibility of said technology. Essentially, this assumption abstracts away from the potential information-generating role of effort. As mentioned in Section 1.1, the literature on experimentation focuses on this

¹¹See (March (1962), Gibbons (2020)) for the notion of “the executive . . . [as] a political broker who cannot solve the problem of conflict by simple payments to the participants and agreement on a superordinate goal.”

role and models an exploitation–exploration trade-off that is technological: exploring a risky option means giving up the myopic gains from exploiting the safe option. Our modeling choice enables us to make the different point that even if information arrives exogenously, which means a single decision-maker would face no trade-off between exploration (competition) and exploitation (collaboration), using this information to inform future choices can still entail costly distortions because of agent incentives.

Continuous time. A finite-horizon continuous-time model allows us to simplify the optimization problem and derive qualitative features of the optimal policy. In particular, in our proofs, working with a Brownian shock process enables us to use time-change methods to cleanly compare selection rules with alternative timing of collaboration.

3. BENCHMARKS

We start with two benchmark settings. First, we characterize a first-best solution, maximizing the principal's ex ante expected profit in the absence of agent incentive constraints. Next, we describe the equilibrium of the three-player game in which the principal cannot commit to a decision rule and must make a static project choice when the deadline arrives.

3.1 First-best solution: Ignoring agent incentives

Toward defining the principal's first-best problem, let $\mathcal{A}$ denote the set of agent strategy profiles $(a_t^1, a_t^{-1})_t$, and let $\mathcal{Y}$ denote the set of project selection rules, i.e., $\{-1, 1\}$ -valued random variables on $\mathcal{F}$. We want to solve the planner problem

$$\begin{aligned} \sup_{a \in \mathcal{A}, y \in \mathcal{Y}} \quad & \mathbb{E} X_T^y = \frac{1}{2} \mathbb{E}[y \Delta X_T] \\ \text{subject to} \quad & dX_t^i = i \Delta a_t \, dt + dB_t^i, \, X_0^1 = x_0^1, \, X_0^{-1} = x_0^{-1}. \end{aligned}$$

The proposition below shows the first-best solution is for the principal to choose the project with the higher output at the deadline and, at every instant before the deadline, have both agents collaborate on the current leader. One part is obvious: the principal will clearly choose the better project ex post. We also show that at any instant before the deadline, it is optimal to have the agents collaborate on the current best guess of which project will be ultimately chosen, so that the effort is productive. Formally, we observe that it is optimal to set $\Delta a_t = 1$ when $\Delta X_t > 0$ and $\Delta a_t = -1$ when $\Delta X_t < 0$.

PROPOSITION 1. *The following policy attains the principal's first-best profit:*

- The principal chooses project $y^{FB} = \ell_T$, the leader as of time T .
- Each agent works on the current leader, that is,

$$(a_t^1, a_t^{-1}) = \begin{cases} (1, 0) : & X_t^1 \geq X_t^{-1} \\ (0, 1) : & X_t^1 < X_t^{-1}. \end{cases}$$

The intuition for this result is straightforward. Because the principal will optimally choose the ex post best project, her objective can be rewritten as $\mathbb{E}|\Delta X_T|/2$, an increasing transformation of $(\Delta X)^2$. However, then the given control increases the drift of $(\Delta X)^2$ more than any other control does, at any given level of $(\Delta X)^2$. A classic comparison theorem from the theory of stochastic differential equations (Ikeda and Watanabe (1977)) says this control yields a stochastically maximal distribution of $(\Delta X_T)^2$.

3.2 No principal commitment

It is immediate that if the principal could not commit, she would (as in the above first-best solution) choose the leading project when the deadline arrives. In other words, the principal's behavior will be ex post optimal: $y = \ell_T$. This observation in turn implies no collaboration will occur, with each agent finding it dominant to devote all his effort to his own project to maximize the chance that it is the eventual winner. Indeed, consider any effort decision of agent $-i$ and any hypothetical effort choice a^i for agent i . Raising a^i to 1 (i.e., never collaborating) increases agent i 's payoff weakly in every state, and strictly with positive probability if he were not already almost surely making the latter choice at almost every time.

PROPOSITION 2. *If the principal cannot commit, any equilibrium is as follows:*

- *The principal chooses project $y^{FB} = \ell_T$, the leader as of time T .*
- *Each agent works on his own project; that is, $(a_t^1, a_t^{-1}) = (1, 1)$.*

3.3 Can commitment be useful?

The first-best solution has the agents collaborating at every instant, whereas in the equilibrium with no principal commitment, getting any collaboration off the ground is not possible. The natural question is whether a principal with some commitment power can foster some collaboration to get better equilibrium outcomes. In this section, we demonstrate informally that if the principal could commit to a decision rule ex ante, she may be able to improve her payoff. Such a principal could, for example, offer any of the following contracts.

Principal's pet project. Consider a project selection rule where the principal simply commits to picking her pet project i , and has agents collaborate on the pet project.

This rule maximizes the benefits of collaboration on the favored project for length of time T , completely forgoing the benefits of choosing the correct project. It is easy to see that this contract can indeed outperform the no-commitment outcome for a range of parameter values.

Unassailable lead. Another project selection rule for the principal has the agents start out competing and commits to picking a project irreversibly if it is the first to take on a lead of at least L . The agents then collaborate on this chosen project. If no project ever reaches the lead threshold L , the principal chooses the leader at time T .

Note that this contract can also improve upon the no-commitment outcome, again by curtailing competition and allowing some collaboration on a favored project. This time, the favored project is the early leader who “wins” the initial competition by overtaking by a specified lead threshold.

Early-lead advantage. Finally, consider the following more elaborate selection rule that gives an early leader an advantage, though the advantage is not unassailable. If no project ever attains a lead L , then the principal chooses the leader at time T . If a project i is the first to take on a lead of L and project $-i$ never catches up, the principal chooses this early leader i with probability $p \geq 1/2$ at the deadline. If project $-i$ subsequently catches up so that agent i 's lead is reduced to 0 and agent $-i$ remains ahead at the deadline, project i is chosen only with probability $2p - 1$. However, if project i is again ahead at the deadline, i is chosen with probability 1. Under this rule, agents start out competing. When an early leader i emerges, both agents start collaborating on the early-leader project and start competing again if and when the early lead disappears.

This selection rule is incentive compatible; in particular, agents willingly collaborate when an early leader emerges because the probability of the early leader being chosen at the deadline, conditional on the lead disappearing, is still $(1/2) + [(2p - 1)/2] = p$. Moreover, such a rule can dominate the outcomes under both the no-commitment and the unassailable-lead contracts.

These examples offer two key takeaways. (i) They demonstrate that commitment power can indeed help the principal improve upon the non-commitment outcome by inducing some collaboration. (ii) Because agent interests are directly opposed, one might have conjectured that the only way to make agent i collaborate is to commit to abandoning project i once and for all. The early-lead-advantage contract demonstrates this reasonable conjecture is actually false: an agent can be willing to collaborate even when he knows that, with sufficient improvement, his project can be chosen in the future.

Indeed, the space of all history-dependent contracts is large and rich, and the substance of our main result is to identify the uniquely optimal one.

4. AGENT INCENTIVES AND THE PRINCIPAL'S PROBLEM

Recall that the principal can choose an arbitrary $\{-1, 1\}$ -valued random variable y on $(\Omega, \mathcal{F}, \mathbb{P})$, and her expected payoff $\mathbb{E}[X_T^y]$ depends on the resulting agent behavior. So we begin by expressing agent incentives more concretely. An agent's strategy is incentive compatible if it maximizes the agent's expected utility (continuation value), given the principal's selection rule. Recall that agent 1's continuation value at time t is q_t and agent -1 's is $-q_t$, where $q_t := \mathbb{E}[y|\mathcal{F}_t]$ describes the interim expected project choice. Notice that $\{q_t\}_t$ is an $\{\mathcal{F}_t\}_t$ martingale. Since $\{\mathcal{F}_t\}_t$ is the filtration generated by the two independent Brownian motions $\Delta X_t - \int_0^t 2\Delta a_s \, d = \Delta B_t$ and $\Sigma X_t = \Sigma B_t$ (see footnote 9), we can apply the martingale representation theorem (Karatzas and Shreve (1998, Theorem 3.4.15)). Specifically, we can represent q as a stochastic integral against these two processes. More precisely, a progressively measurable $\mathbb{R}^2$ -valued process $\{C_t = (c_t^\Delta, c_t^\Sigma)\}_t$

on filtration $\{\mathcal{F}_t\}_t$ exists whose time- t quadratic variation has finite expectation for every $t \geq 0$ and such that¹²

$$dq_t = [c_t^\Delta (d\Delta X_t - 2\widetilde{\Delta a}_t dt) + c_t^\Sigma d\Sigma X_t], \quad (\text{LoM})$$

where $\widetilde{\Delta a}_t$ is the equilibrium-anticipated Δa_t , and the law of motion $d\Delta X_t$ is influenced by the chosen Δa_t . Intuitively, we can think of c_t^Δ and c_t^Σ as project sensitivities that describe how the interim expected project choice responds to relative and aggregate shocks of the two projects. Because agents' contemporaneous choices influence $d\Delta X_t$ but not $d\Sigma X_t$, agent incentive compatibility requires

$$\Delta a_t = 0 \quad \text{whenever } c_t^\Delta \neq 0. \quad (\text{IC})$$

Indeed, given $c_t^\Delta > 0$ (resp. $c_t^\Delta < 0$), both agents would have a strict incentive to choose $a_t^i = 1$ (resp. $a_t^i = 0$). Further, we can rewrite the principal's profit, Π , as

$$\Pi = \frac{1}{2} \mathbb{E}[y\Delta X_T] = \frac{1}{2} \mathbb{E}[q_T \Delta X_T].$$

Applying Ito's formula on the product process $q \cdot \Delta X$ and using the independence of the two processes ΔB and ΣB , direct computation shows

$$\begin{aligned} \Pi &= \frac{1}{2} q_0 \Delta X_0 + \frac{1}{2} \mathbb{E} \left[\int_0^T (q_t d\Delta X_t + \Delta X_t dq_t) \right] + \mathbb{E} \left[\int_0^T dq_t d\Delta X_t \right] \\ &= \underbrace{\frac{1}{2} q_0 \Delta X_0}_{\text{ex ante}} + \mathbb{E} \int_0^T \left(\underbrace{q_t \Delta a_t}_{\text{collaboration}} + \underbrace{c_t^\Delta}_{\text{adaptivity}} \right) dt, \end{aligned}$$

where the last equality comes from the standard formula for quadratic covariation of stochastic integrals. Writing down the objective in this way clarifies that, apart from making an appropriate ex ante project choice (choosing q_0), the principal has two levers to increase profit: (i) adapting the project choice to relative productivity shocks, which will (by (IC)) induce agents to compete on their respective projects; and (ii) eliminating the efficiency loss of competition by having agents collaborate.

Whenever c_t^Δ is 0 (and so $\Delta a_t \in [-1, 1]$ is not restricted by agent incentives), changing Δa_t to $\text{sign } q_t \in \{-1, 0, 1\}$ raises the principal's objective.¹³ Moreover, this change does not alter feasibility of (q, c^Δ, c^Σ) . Therefore, we can solve out this choice variable and write the principal's objective as

$$\Pi = \frac{1}{2} q_0 \Delta X_0 + \mathbb{E} \int_0^T (\mathbb{1}_{c_t^\Delta=0} |q_t| + c_t^\Delta) dt.$$

Recall that q_t is the interim expected project choice. The choice of Δa_t being 1 (resp. -1) when q_t is positive (resp. negative) corresponds to the principal choosing collaboration

¹²That is, $q_t = q_0 + \int_0^t [c_s^\Delta (d\Delta X_s - 2\widetilde{\Delta a}_s ds) + c_s^\Sigma d\Sigma X_s]$ almost surely.

¹³Moreover, it strictly raises the objective if, with positive probability for a positive measure of times t , we have $q_t \neq 0$ and $\Delta a_t \neq 0$.

on the “current favorite” at time t : the project with a higher likelihood of being chosen eventually given history up to t . Note that the current favorite may not be the project that is currently ahead, because the former is endogenous to the principal’s chosen contract. Determining what the current favorite is at any time is an essential part of the characterization of the optimal contract. We can write the principal’s problem as

$$\sup_{(q_t, c_t^\Delta, c_t^\Sigma)_t} \left\{ \frac{1}{2} q_0 \Delta X_0 + \mathbb{E} \int_0^T (\mathbb{1}_{c_t^\Delta=0} |q_t| + c_t^\Delta) dt \right\}, \quad (\text{O})$$

$$\text{subject to } q_t \in [-1, 1] \text{ and } dq_t = c_t^\Delta d\Delta B_t + c_t^\Sigma d\Sigma B_t. \quad (\text{P})$$

5. THE OPTIMAL SELECTION RULE

Our main result describes the form of the uniquely optimal project selection rule.

THEOREM 1. *An optimal contract exists and is unique.¹⁴ A bounded, continuous, non-decreasing function $\bar{z} : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ with $z_0 = 0$ and $z_t > 0$ for every $t > 0$ exists, such that (whatever is the duration T until the deadline) the following behavior is optimal:*

- *The principal chooses project $y^* = \ell_{\tau^*}$, the leader as of time¹⁵*

$$\tau^* := \inf\{t \in [0, T] : |\Delta X_t| \geq \bar{z}_{T-t}\}.$$

- *Each agent works on his own project before τ^* .*
- *Both agents work on project y^* from time τ^* onward.*

Figures 1 and 2 show realized paths of projects’ relative performance in which projects 1 and -1 are chosen, respectively. Consistent with the theorem, a winner is chosen and permanent collaboration begins the first time one project’s lead exceeds the threshold. Figure 3 demonstrates a realization in which the project choice (project 1) turns out to be ex post inefficient.

Recall that the theorem is stated with the normalizations $\sigma = \mu = 1$, $\beta = -1$, and $\Sigma X_0 = 0$. However, it is easy to deduce the optimal contract for general parameter values. In particular, the initial aggregate state (ΣX_0) and baseline development rate (β) are irrelevant, for the same reasons that aggregate performance is optimally ignored. Because each quantity enters the threshold rule in two places, comparative statics with respect to the marginal product of effort (μ) and the project volatility (σ) are more delicate, depending on detailed features of the function $\bar{z}$. However, some features are easy; for instance, raising μ while holding the ratio μ/σ fixed raises the lead threshold at every time. Finally, extending the deadline (T) leads to a larger lead threshold at every time, whereas changing the initial project asymmetry (ΔX_0) in favor of one project has no effect on the standards and, hence, makes that project more likely to be chosen.

¹⁴Any two optimal incentive-compatible selection rules almost surely have the same chosen project and the same agent choices at almost every time.

¹⁵In the zero-probability event that $\tau^* = T$ and $X_T^1 = X_T^{-1}$, the principal may choose arbitrarily.




FIGURE 1. Project 1 is selected.




FIGURE 2. Project -1 is selected.

The remainder of the section is dedicated to proving the theorem. A common approach to solving a stochastic control problem like the one in (O) is to heuristically derive the HJB equation, establish the existence of a smooth solution to it, and appeal to a verification theorem. However, this approach has two limitations. First, given a finite horizon, the HJB equation would be a partial differential equation, and so explicit characterization of its solutions is not straightforward. Second, note that (O) is a volatility control problem. Discontinuous flow payoffs and the possibility of degenerate volatility (which happens when c^Σ and c^Δ are 0) mean that the HJB equation need not be uniformly elliptical. Moreover, the associated stochastic differential equation (SDE) for




FIGURE 3. An ex post inefficient choice.

q_t in (P) may not have a strong solution.¹⁶ So we adopt a different route. It turns out that time-change methods for analyzing weak controls can simplify our volatility control problem. Because the argument is not typical of the optimal contracting literature, we first summarize our approach and then proceed to formal analysis.

5.1 Approach to characterizing the optimal selection rule

Our first technical step is to consider relaxations of the principal's problem that allow for *weak* solutions. Recall that in the control problem (O), the principal has to choose interim expected project choice q and project sensitivities C . Allowing weak solutions means we now allow the principal to additionally choose the underlying Brownian motions that drive projects' random evolution (while still respecting the *law* governing this evolution as stated in the model section).¹⁷ In a typical discrete-time model, such a relaxation would be irrelevant, but in the present setting, it is a useful tool for the analyst. Given this broader definition of a control, we then proceed to show that restricting attention to controls that have various economically intuitive features is without loss of optimality.

The first step toward establishing the result is Lemma 1. In a quantitative argument, we establish that $\|(c^\Delta, c^\Sigma)\| \geq 1$. Intuitively, the principal can back-load collaboration by continuously speeding up decision-making whenever $\|(c^\Delta, c^\Sigma)\|$ is too small, thus creating residual time at the end for more collaboration.

Using this, in Lemma 2, we show that the principal ignores aggregate shocks. To this end, notice that c^Σ does not affect the principal's objective function or the agent

¹⁶We thank an anonymous referee for crystallizing these specific issues with the HJB approach.

¹⁷Other papers on continuous-time optimal contracting and elsewhere in the optimal control literature appeal to weak solutions of stochastic differential equations (e.g., Sannikov (2008)). We believe our specific use of weak solutions is novel. Namely, we solve a relaxed program entailing weak solutions as a solution method to characterize principal-optimal incentive-compatible (IC) strong solutions.

incentives, and that aggregate shocks are not an informative signal of agents' choices. Therefore, it is intuitive that the principal should set $c^\Sigma = 0$. A consequence of these two lemmas is that $|c^\Delta| \geq 1$ until $q_t \in \{-1, 1\}$. That is, if an optimal contract exists, it will back-load all the collaboration ($c^\Delta = 0$) and front-load all the competition.

In Lemma 3, we show that we can reduce the principal's problem to an optimal stopping problem in which the principal chooses a time when she stops competition and switches the agents to collaboration on the current leader as of that moment, until the deadline. In particular, it is optimal for the principal to have the agents stop competing at some time, make a constrained-efficient choice with the partial information she has, and switch to collaboration on the chosen project thereafter.

Once we have established competition is ended once and for all, showing that the threshold for doing so should decline over time is straightforward. We formalize this fact in Lemma 4 by showing that the stopping rule is a decreasing threshold: the principal switches to collaboration on a project as soon as its lead over the other project is sufficiently large, with this lead standard becoming less demanding as the deadline approaches.¹⁸

In the final step, we show that even though the above qualitative features are derived for relaxations of the principal's problem, these relaxations are payoff-irrelevant: The constructed optimal control from the relaxed problem can be implemented in the original problem through strong solutions, i.e., with the principal choosing simply q and C as in the original problem.

5.2 Mathematical preliminaries

We start by defining a permissive notion of a weak control that will be convenient.

DEFINITION 1. A *weak control* is a tuple $\mathcal{C} = \langle \Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P}, B, C, q \rangle$ such that the following definitions hold:

- (i) The term $\langle \Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P} \rangle$ is a filtered probability space satisfying the usual conditions.
- (ii) The term $B = (\Delta B, \Sigma B) = \{B_t\}_{t \geq 0}$ is an $\mathbb{R}^2$ -valued stochastic process on its natural filtration $\{\mathcal{F}_t\}_{t \geq 0}$, such that $B/\sqrt{2}$ is a standard Brownian motion.
- (iii) The term $C = (c^\Delta, c^\Sigma) = \{C_t\}_{t \geq 0}$ is a progressively measurable $\mathbb{R}^2$ -valued process on $\{\mathcal{F}_t\}_{t \geq 0}$ whose time- t quadratic variation has finite expectation for every $t \geq 0$.
- (iv) The term $q = \{q_t\}_{t \geq 0}$ is a $[-1, 1]$ -valued martingale on $\{\mathcal{F}_t\}_{t \geq 0}$.
- (v) The term $q_t = q_0 + \int_0^t C \cdot dB$ almost surely while $|q_t| < 1$.

Defining the notion of a Brownian base is also convenient. A *Brownian base* is any tuple $\langle \Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P}, B \rangle$ satisfying properties (i) and (ii) above.

¹⁸Studying costly sequential sampling problems that a single decision-maker faces, Fudenberg, Strack, and Strzalecki (2018) show that decreasing threshold rules can arise even without a deadline.

Note that if a principal chooses a weak control as defined above, she also chooses the underlying stochastic process and probability space. Of course, in our principal's problem in (O), she has no such choice. She must take a particular Brownian base $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P}, B)$ as given. However, considering this relaxation of the principal's problem is convenient.

Given a weak control $\mathcal{C} = (\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P}, B, C, q)$, we define

$$\tau_{\mathcal{C}} := T \wedge \inf\{t \in [0, T) : |q_t| = 1\},$$

$$J(\mathcal{C}) := \frac{1}{2}q_0\Delta X_0 + \mathbb{E}\left[\int_0^{\tau_{\mathcal{C}}} (\mathbb{1}_{c_t^{\Delta}=0}|q_t| + c_t^{\Delta}) \, dt + T - \tau_{\mathcal{C}}\right].$$

Intuitively, given a weak control, $\tau_{\mathcal{C}}$ is the stopping time associated with that control when q_t hits a boundary, or when the principal has no choice left to make, and $J(\mathcal{C})$ is the payoff that the principal would get if the control were followed until $\tau_{\mathcal{C}}$ and then agents collaborated on the choice at $\tau_{\mathcal{C}}$.

5.3 Decide quickly

We first establish a quantitative claim about optimal selection rules. For the principal to resolve uncertainty somewhat quickly is without loss of optimality. Specifically, if the principal is deciding which project to choose slowly enough that its flow benefits are smaller than those from collaboration on a chosen project, she can improve her payoff by speeding up her decision-making (higher $\|C\|$) and defer any saved time toward end-game collaboration on her chosen project. Formally, this claim amounts to showing that restricting attention to weak controls such that $\|\hat{C}\| \geq 1$ is without loss of optimality.

LEMMA 1. *For any weak control $\mathcal{C}$, a weak control $\hat{\mathcal{C}}$ exists whose Euclidean norm satisfies $\|\hat{C}\| \geq 1$ and such that $J(\hat{\mathcal{C}}) \geq J(\mathcal{C})$. Moreover, $J(\hat{\mathcal{C}}) > J(\mathcal{C})$ unless, almost surely, $\|C_t\| \geq 1$ for almost every $t \in [0, T)$ with $|q_t| < 1$.*

In addition to telling us uncertainty is resolved quickly, Lemma 1 is a key ingredient to Lemma 2, which proves that $c^{\Sigma} = 0$, i.e., the principal does not update her choice in response to aggregate performance of the two projects. The interested reader can refer to the [Appendix](#) for the proof, but we summarize the logic here.

The proof of Lemma 1 is constructive, modifying a weak control without this property to a superior one with this property. Specifically, the fractal property of Brownian motion allows us to construct a superior weak control by replacing the underlying Brownian motion with a law-equivalent time change of the same, and replacing our sensitivity coefficient C with one that is scaled up whenever the original one had $\|C\| < 1$, in such a way that the expected project choice q follows the same trajectory. Intuitively, this argument is akin to “slowing down the clock” without changing the trajectory of the expected project choice, thus simply speeding up the original decision-making process and creating some residual time at the end. The benefit of rescaling time in this way is that this “extra” residual time can be utilized for efficient collaboration on a chosen

project for a flow benefit of 1. Of course, the cost of this speeding up is that the duration for collecting flow payoffs is reduced. Note that holding an expected project choice q_0 fixed, the principal's payoff in (O) can be interpreted as the sum of the total accrued value of collaboration ($\int_0^T \mathbb{1}_{c_t^\Delta=0} |q_t| dt$) and the total net value of competition ($\int_0^T c_t^\Delta dt$), since agents compete whenever $c_t^\Delta \neq 0$. Thus, the foregone flow payoff is either c_t^Δ (if from competition) or $\mathbb{1}_{c_t^\Delta=0} |q_t|$ (if from collaboration), both of which are bounded above by 1. Thus, the cost of lost flow payoff as a result of speeding up decision-making is always less than the benefit of the extra collaboration time.

5.4 Respond only to relative performance

We next establish that for the principal to respond to the relative performance of projects, and not to aggregate shocks, is without loss of optimality. Absent an agency problem, such a choice is, of course, allocatively efficient; we show that this property remains optimal even when respecting agent incentives.

LEMMA 2. *For any weak control C , a weak control $\hat{C}$ exists that satisfies $\hat{c}^\Delta \geq 1$ and $\hat{c}^\Sigma = 0$, and such that $J(\hat{C}) \geq J(C)$. Moreover, $J(\hat{C}) > J(C)$ unless, almost surely, $c^\Delta \geq 1$ and $c^\Sigma = 0$ for almost every $t \in [0, T)$ with $|q_t| < 1$.*

We show constructively that restricting attention to weak controls that ignore aggregate shocks and respond to relative shocks (i.e., set $\hat{c}^\Sigma = 0$ in such a way that $\|\hat{C}\| = \|C\|$, which leaves $\hat{c}^\Delta \geq 1 > 0$) is without loss of optimality. By responding solely to contemporaneous relative shocks while maintaining the degree to which she resolves uncertainty based on current shocks, the principal can better capitalize on the gains of competition today while keeping the law of q_t fixed, and so without affecting her ability to respond optimally in the future. Such a change will still entail a potential cost of foregone current collaboration, but if the principal is resolving uncertainty sufficiently quickly (which she optimally does by Lemma 1), these costs are smaller than the gains to more effective competition.

For the interested reader, the proof of this lemma is a good example of why weak solutions are especially useful. The conclusion of Lemma 2—that $\hat{c}^\Sigma$ is almost surely 0—is a natural conjecture because c_t^Σ neither affects players' incentives nor enters the objective function. However, establishing this conjecture formally (by standard methods) would require us to write down the HJB equation corresponding to (O) and prove both the existence of a solution and concavity. The difficulty is that with a discontinuous flow payoff, our HJB equation does not belong to a class of well understood partial differential equations. Reasoning through weak solutions circumvents this challenge. To see the basic idea, suppose that, on some paths, we have $c_t^\Sigma \geq 1$ and $c_t^\Delta = 0$. Then we would conjecture that swapping c_t^Σ and c_t^Δ on those histories would lead to a payoff improvement (flow payoff being $|q_t| \leq 1 \leq c_t^\Sigma$) while keeping the law unchanged. However, such a swapping operation changes the path of q_t . Permitting weak solutions enables us to construct new Brownian motions using the original ones and to work with these alternate weak controls that deliver a payoff improvement while preserving the law of

motion. Although swapping arguments of this kind have been commonly used in studying dynamic incentive problems in discrete time, we hope our approach can facilitate similar arguments in continuous-time models.

5.5 First compete and then collaborate on the winner

Next we show we can bound the payoff attainable in the present optimal control problem by an optimal stopping problem. The principal's problem reduces to one in which she picks a stopping time at which she switches from pure competition to permanent collaboration on the chosen project until the deadline.

LEMMA 3. *For any weak control $\mathcal{C}$ such that $c^\Delta \geq 1$ and $c^\Sigma = 0$, the stopping time $\tau := \tau_{\mathcal{C}}$ has $J(\mathcal{C}) \leq T + \mathbb{E}[(1/2)|\Delta X_0 + \Delta B_\tau| - \tau]$. Moreover, $J(\mathcal{C}) < T + \mathbb{E}[(1/2)|\Delta X_0 + \Delta B_\tau| - \tau]$ unless, almost surely, q_τ is equal to the sign of $\Delta X_0 + \Delta B_\tau$ if $\Delta X_0 + \Delta B_\tau \neq 0$.*

The lemma follows from a direct computation of $J(\mathcal{C})$, given properties of $\mathcal{C}$. The details are provided in the [Appendix](#). This lemma reduces our principal's problem to a nonstrategic optimal stopping problem. To see why, notice that $\{Y_t = (\Delta X_0 + \Delta B_t)/2\}_t$ is an exogenous Brownian motion. Subtracting the constant term T from the principal's objective, the lemma says her modified objective is no greater than

$$\sup_{\text{stopping times } \tau \leq T} \mathbb{E}[|Y_\tau| - \tau],$$

which is exactly the value of an optimal stopping problem with constant flow cost, terminal value $|Y|$, and deadline T . As with our other lemmas, we eventually show that this upper bound is attained, and so the two problems are equivalent.

5.6 The winner is the first to take a large enough lead

The final building block is to show the optimal stopping rule takes an intuitive form, permanently switching to collaboration on the leading project when it first takes a large enough lead, where the standard for “large enough” becomes less demanding as the deadline approaches.

LEMMA 4. *A function $\bar{z} : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ exists, such that for any Brownian base $\mathcal{B}$ and any $(T, z) \in \mathbb{R}_+ \times \mathbb{R}$, $\tau_{T,z,\mathcal{B}}^* := \inf\{t \in [0, T] : |z + \Delta B_t| \geq \bar{z}_{T-t}\}$ is a $(\mathcal{B}, T)$ stopping time,¹⁹ and every Brownian base $\hat{\mathcal{B}}$ and $(\hat{\mathcal{B}}, T)$ stopping time $\hat{\tau}$ have*

$$\mathbb{E}[(1/2)|z + \Delta B_{\tau_{T,z,\mathcal{B}}^*}| - \tau_{T,z,\mathcal{B}}^*] \geq \mathbb{E}[(1/2)|z + \widehat{\Delta B}_{\hat{\tau}}| - \hat{\tau}],$$

with equality if and only if $\hat{\tau}$ is almost surely equal to $\tau_{T,z,\hat{\mathcal{B}}}^*$. Moreover, $\bar{z}$ is bounded, continuous, and nondecreasing, with $\bar{z}_0 = 0$ and $\bar{z}_T > 0$ for every $T > 0$.

¹⁹A $(\mathcal{B}, T)$ stopping time is a stopping time on the filtration underlying $\mathcal{B}$ that respects deadline T . See Definition 2 in the [Appendix](#).

The proof is provided in the [Appendix](#), but we sketch the argument here. We have an optimal stopping problem in which a decision-maker observes a driftless Brownian motion at a constant flow cost and can stop at any time before a deadline, where stopping yields a payoff equal to the Brownian motion's absolute value. The finite deadline makes the problem nonstationary, and so we do not attempt to derive a closed-form solution for the optimal stopping rule, but instead derive qualitative features of it. Classic results from the optimal stopping literature imply that in our problem, the earliest optimal policy is to stop as soon as the optimal and stopping values coincide, and we can, in fact, show that this stopping time is uniquely optimal. Thus we analyze the (continuous) optimal value function, taking as arguments the time remaining and the current state of the Brownian motion, and show that the set of values of the Brownian motion at which the optimal value function strictly exceeds the absolute value (stopping value) is a bounded, symmetric, nonempty interval that shrinks as the deadline approaches. Boundedness obtains by considering a relaxed problem with no deadline and using existing results for problems with an infinite horizon. The set shrinks as the deadline approaches, because the decision-maker's objective is unchanged but is subject to a tighter constraint. A limit argument shows that it contains 0 when near enough to the deadline, and, hence (given monotonicity), contains 0 at every time. It is symmetric about 0 because the objective and the law of motion are. Finally, it is an interval around 0 because the value function is convex, whereas the terminal value is affine on either side of 0.

5.7 Characterizing the optimal selection rule

We next show that the qualitative features we derived for the solution to the relaxed principal's problem (permitting weak solutions) also apply to the optimal rule in our original problem. Accordingly, the unique optimal selection rule takes the simple form described in our main theorem.

PROOF OF THEOREM 1. Taking $\bar{z}$ to be the function delivered by Lemma 4, let Π^* be the principal value generated by the behavior named in the theorem. Note that the described agent behavior is incentive compatible given this selection rule: agents are indifferent from τ^* until the deadline, and they increase their probability of being the time- τ^* leader by working on their own projects.

Consider now an arbitrary selection rule by the principal, together with incentive-compatible agent behavior, and let Π be the principal's value from adopting it. As we have shown in Section 4, it generates some weak control $\mathcal{C}$ such that $J(\mathcal{C}) = \Pi$.

Now let us apply the lemmas referenced above. Lemma 2 delivers some weak control $\hat{\mathcal{C}}$ such that $\hat{c}^\Delta > \hat{c}^\Sigma = 0$ and such that $J(\hat{\mathcal{C}}) \geq \Pi$, the latter inequality being strict unless, almost surely, $c^\Delta > c^\Sigma = 0$ for almost every $t \in [0, T)$ with $|q_t| < 1$ (in which case, we can take $\hat{\mathcal{C}} = \mathcal{C}$ without loss). Lemma 3 then tells us the stopping time $\hat{\tau} := \tau_{\hat{\mathcal{C}}}$ has $T + \mathbb{E}[(1/2)|\Delta X_0 + \widehat{\Delta B}_{\hat{\tau}}| - \hat{\tau}] \geq \Pi$, strictly so unless $q_{\hat{\tau}}$ is almost surely equal to the sign of $\Delta X_0 + \widehat{\Delta B}_{\hat{\tau}}$ if $\Delta X_0 + \widehat{\Delta B}_{\hat{\tau}} \neq 0$. Finally, Lemma 4 tells us that τ^* (as defined in the statement of the theorem) has $T + \mathbb{E}[(1/2)|\Delta X_0 + \Delta B_{\tau^*}| - \tau^*] \geq \Pi$, strictly so unless $\hat{\tau}$ is almost surely equal to $\tau_{T,z,\hat{\mathcal{C}}}^*$.

The above arguments directly deliver the theorem. First, they show that the principal's optimal value is $\Pi^* = T + \mathbb{E}[(1/2)|\Delta X_0 + \Delta B_{\tau^*}| - \tau^*]$, making the described behavior principal-optimal. Second, they establish that $\Pi < \Pi^*$ (making the given selection rule and agent behavior suboptimal) unless, almost surely, the selected project is the same and agent choices are the same at almost every time. $\square$

6. DISCUSSION

6.1 *Duration of collaboration and ex post inefficiency*

An implication of our characterization of the optimal contract is that the length of the competition phase is probabilistically bounded in two senses. First, for any deadline T , a phase of collaboration always exists, because the threshold collapses as the deadline approaches. Second, if we increased the time horizon T , although the duration of the competition phase would increase (in the sense of first-order stochastic dominance), the duration of competition would remain uniformly bounded.²⁰ Put differently, not only is the collaboration phase reached with probability 1 for any T , but also, when the project is of a very long-term nature, most of its development is spent collaborating.

Fostering collaboration increases the value of the principal's chosen project, but the inefficiency caused by picking the “wrong” project on-path can be arbitrarily large; that is, given any $M > 0$, the probability that $X_T^y + M < X_T^{-y}$ is strictly positive.²¹ Nevertheless, because collaboration starts early in expectation, the probability of an error approaches 0 as the project horizon grows long.

6.2 *Cancellation of projects before the deadline*

In our setting, the principal chooses an optimal stopping time at which she makes a choice and then has both agents collaborate on the chosen project. An alternative interpretation is one in which the principal chooses when to irreversibly cancel one of the projects, after which both agents must work on the remaining project. A richer contracting environment in which the principal can choose to irreversibly cancel is arguably more consistent with our motivating application and is indeed equivalent to our current model. On the one hand, the principal cannot be worse off in the richer environment, because she can always abstain from cancelling projects. Conversely, the principal can always simulate cancellation through a selection rule by deciding on a project in advance and having the agents collaborate on the chosen project. Moreover, in the richer model allowing irreversible termination, our optimal selection rule from Theorem 1 can

²⁰That is, some finite-mean random variable τ_∞ exists such that the duration of the competition phase is first-order-stochastically dominated by τ_∞ . Indeed, one could take τ_∞ to be the optimal stopping time from an analogous stopping problem with no deadline, which is known to exhibit a constant lead threshold: the proof of Lemma 4 notes that it is $\frac{1}{\sqrt{2}}$. The constant threshold is finite because the option value of continuing vanishes with the probability of the Brownian motion revisiting zero, and a finite $|\Delta X|$ threshold is surpassed in finite expected time because the constant volatility is nonzero.

²¹This observation is a consequence of the projects' evolution being a Brownian motion and, thus, having unbounded supports.

be implemented in equilibrium, *without commitment*. The principal could simply terminate the project that is lagging behind by the current lead threshold, with each agent working on his own project unless it is cancelled.

6.3 Agent indifference in the collaboration phase

Under our optimal rule, when agents collaborate, they are indifferent between competing and collaborating. Such indifference is common in many standard contracting environments, but constructing a similar contract with strict incentives is typically possible using a small monetary perturbation.²² In our setting with no transfers, no obvious way exists to turn weak incentives into strict ones. Indeed, the best that the principal can achieve in any strict equilibrium is simply the no-commitment solution.²³

However, following the spirit of classical contract theory arguments, we ask what would happen if the principal could provide monetary incentives up to a fixed budget of $\epsilon > 0$ (assuming for simplicity that agents have separable preferences over money and project choice). Observe that, in this case, implementing the optimum from Theorem 1 in a strict equilibrium would be easy.²⁴ For example, letting τ^* and y^* be as defined in the statement of Theorem 1, the principal could augment the given selection rule by further giving both agents a prize of $\epsilon/2$ if and only if the collaboration phase goes better than expected, that is,

$$X_T^{y^*} > \mathbb{E}_{\tau^*}[X_T^{y^*}] = X_{\tau^*}^{y^*} + (\beta + 2\mu)(T - \tau^*).$$

This contract clearly gives agents strict incentives to collaborate on the chosen project, and it does not distort incentives in the initial phase of competition, because an agent's monetary prize of $\epsilon/2$ will be earned with probability 1/2 conditional on any outcome of the initial competition. Moreover, one can show that the principal's optimal value converges to ours as $\epsilon \rightarrow 0$. One can also similarly modify the model with small-scale monetary incentives to accommodate a small effort cost or a small preference for an agent to work on his own project.

Hence, our model is perhaps best interpreted as a parsimonious version of the $\epsilon \approx 0$ model (in which strict incentives are without loss), wherein agents' empire-building motives overwhelm monetary incentives of a realistic scale. The contribution of this paper is to show that, somewhat surprisingly, fostering a degree of collaboration in equilibrium is still possible and optimal despite the paucity of powerful incentivizing instruments.

²²For example, optimally inducing high effort in a textbook binary-action moral hazard model with contractible transfers will leave the agent indifferent, but modifying the contract to provide slightly higher-powered incentives will approximate the same principal value under strict incentive compatibility.

²³Given that agents' interests are directly opposed, every strict equilibrium has the agents always working on opposite projects. The best the principal can do subject to this constraint, then, is to choose the best project ex post.

²⁴We focus here on perturbing the model by allowing small-scale monetary transfers, but other perturbations would similarly enable strict incentives. For example, essentially identical reasoning would apply if the principal could instead choose to adopt neither project with some small probability.

APPENDIX: OMITTED PROOFS

In this appendix, we provide proofs that we omitted from the main text of the paper.

A.1 *Proof of Proposition 1*

First, because the ex post efficient rule $y = \ell_T$ maximizes the principal's objective state-wise, we may recast her problem as

$$\begin{aligned} \sup_{a \in \mathcal{A}} \quad & \frac{1}{2} \mathbb{E} |\Delta X_T| \\ \text{subject to} \quad & dX_t^i = (a_t^i - a_t^{-i}) dt + dB_t^i, \quad X_0^1, X_0^{-1}. \end{aligned}$$

Now, before showing that the described agent behavior is optimal, observe that our posited optimal weak control is indeed well defined: Following Example 1.2 of [Yamada \(1973\)](#), the stochastic differential equation

$$d\Delta X_t = 2 \operatorname{sign}(\Delta X_t) dt + d\Delta B_t$$

admits a unique strong solution. Optimality then follows readily from a comparison theorem. Indeed, following identically the proof of Theorem 2.1 in [Ikeda and Watanabe \(1977\)](#), any alternative weak control has a (weakly) first-order-stochastically dominated distribution of $|\Delta X_T|$.²⁵ $\square$

A.2 *Proof of Lemma 1*

Let $\tau := \tau_C$, and assume without loss that $c_t^\Delta = 1$ and $c_t^\Sigma = 0$ whenever $t \geq \tau$. Moreover, assume without loss (changing C on a measure zero set) that C is 0 on any time interval where it is a.e. 0.

We now proceed to define our candidate $\hat{C}$. Define

$$\begin{aligned} \gamma_t &:= 1 \wedge \|C_t\| \text{ (where } \|\cdot\| \text{ is the Euclidean norm on } \mathbb{R}^2), \\ \zeta_t &:= \int_0^t \gamma_s^2 ds \text{ (nondecreasing and 1-Lipschitz, with slope 1 after } \tau), \\ \lambda_u &:= \inf\{t \geq 0 : \zeta_t > u\}, \\ \hat{\mathcal{F}}_u &:= \mathcal{F}_{\lambda_u} = \{E \in \mathcal{F}_\infty : E \cap \{\lambda(u) \leq t\} \in \mathcal{F}_t \forall t \geq 0\}, \\ \hat{B}_u &:= \int_0^{\lambda_u} \gamma_t dB_t, \\ \hat{C}_u &:= \begin{cases} \frac{1}{\gamma_{\lambda_u}} C_{\lambda_u} : & C_{\lambda_u} \neq (0, 0), \\ (1, 0) : & C_{\lambda_u} = (0, 0), \end{cases} \end{aligned}$$

²⁵That result shows that a control $\Delta a_t = -\operatorname{sign}(\Delta X_t)$ minimizes $|\Delta X_T|$ —in fact, minimizes each of $\{|\Delta X_t|\}_{t \in [0, T]}$ —in a first-order stochastic differential (FOSD) sense. However, reproducing the proof nearly verbatim establishes that control $\Delta a_t = \operatorname{sign}(\Delta X_t)$ maximizes $|\Delta X_T|$.

$$\hat{q}_u := q_{\lambda_u},$$

$$\hat{C} := \langle \Omega, \mathcal{F}, \{\hat{\mathcal{F}}_u\}_{u \geq 0}, \mathbb{P}, \hat{B}, \hat{C}, \hat{q} \rangle.$$

First, we observe that λ_u is a $\{\mathcal{F}_t\}_{t \geq 0}$ stopping time for each $u \geq 0$ and that the tuple $\langle \Omega, \mathcal{F}, \{\hat{\mathcal{F}}_u\}_{u \geq 0}, \mathbb{P}, \hat{B} \rangle$ is a Brownian base. These facts follow directly from applying the Dambis–Dubins–Schwarz theorem (Karatzas and Shreve (1998, Theorem 3.4.6)) to $M = \frac{1}{\sqrt{2}}\hat{B}$, with the observation that (applying the formula for quadratic variation of an Itô process)

$$\langle M \rangle_t = \int_0^{\lambda_t} \gamma_s^2 \, ds = \zeta_{\lambda_t} = t.$$

To see that $\hat{C}$ is a weak control, all that remains is to check that $\hat{q}_u = \hat{q}_0 + \int_0^u \hat{C} \cdot d\hat{B}$ or, equivalently, that $\int_0^u \hat{C} \cdot d\hat{B} = \int_0^{\lambda_u} C \cdot dB$. This formula follows directly from Proposition 3.4.8 of Karatzas and Shreve (1998).

Let us now show that $J(\hat{C}) \geq J(C)$. To this end, first observe that

$$\begin{aligned} \int_0^{\zeta_\tau} (\mathbb{1}_{\hat{c}_u^\Delta=0} |\hat{q}_u| + \hat{c}_u^\Delta) \, du &= \int_0^\tau (\mathbb{1}_{\hat{c}_{\zeta_t}^\Delta=0} |\hat{q}_{\zeta_t}| + \hat{c}_{\zeta_t}^\Delta) \, d\zeta_t \\ &= \int_0^\tau (\mathbb{1}_{\hat{c}_{\zeta_t}^\Delta=0} |\hat{q}_{\zeta_t}| + \hat{c}_{\zeta_t}^\Delta) \gamma_t^2 \, dt \\ &= \int_0^\tau (\mathbb{1}_{c_t^\Delta=0} |q_t| \gamma_t + c_t^\Delta) \gamma_t \, dt, \end{aligned}$$

so that

$$\begin{aligned} \tau - \zeta_\tau + \int_0^{\zeta_\tau} (\mathbb{1}_{\hat{c}_u^\Delta=0} |\hat{q}_u| + \hat{c}_u^\Delta) \, du - \int_0^\tau (\mathbb{1}_{c_t^\Delta=0} |q_t| + c_t^\Delta) \, dt \\ &= \int_0^\tau [1 - \gamma_t^2 + (\mathbb{1}_{c_t^\Delta=0} |q_t| \gamma_t + c_t^\Delta) \gamma_t - (\mathbb{1}_{c_t^\Delta=0} |q_t| + c_t^\Delta)] \, dt \\ &= \int_0^\tau [(1 - \gamma_t^2) - (1 - \gamma_t^2) \mathbb{1}_{c_t^\Delta=0} |q_t| - (1 - \gamma_t) c_t^\Delta] \, dt \\ &= \int_0^\tau (1 - \gamma_t) [(1 + \gamma_t)(1 - \mathbb{1}_{c_t^\Delta=0} |q_t|) - c_t^\Delta] \, dt. \end{aligned}$$

Finally, that $\hat{c}_t^\Delta = 1$ for every $t \geq \zeta_\tau$ implies

$$\begin{aligned} J(\hat{C}) - J(C) &= \mathbb{E} \left[\int_0^{\zeta_\tau} (\mathbb{1}_{\hat{c}_u^\Delta=0} |\hat{q}_u| + \hat{c}_u^\Delta) \, du + T - \zeta_t \right] - \mathbb{E} \left[\int_0^\tau (\mathbb{1}_{c_t^\Delta=0} |q_t| + c_t^\Delta) \, dt + T - \tau \right] \\ &= \mathbb{E} \left[\tau - \zeta_t + \int_0^{\zeta_\tau} (\mathbb{1}_{\hat{c}_u^\Delta=0} |\hat{q}_u| + \hat{c}_u^\Delta) \, du - \int_0^\tau (\mathbb{1}_{c_t^\Delta=0} |q_t| + c_t^\Delta) \, dt \right] \\ &= \mathbb{E} \int_0^\tau (1 - \gamma_t) [(1 + \gamma_t)(1 - \mathbb{1}_{c_t^\Delta=0} |q_t|) - c_t^\Delta] \, dt. \end{aligned}$$

The value ranking will then follow if we establish that $(1 - \gamma_t)[(1 + \gamma_t)(1 - \mathbb{1}_{c_t^\Delta=0}|q_t|) - c_t^\Delta]$ is nonnegative for any $t \in [0, \tau]$, and is strictly positive if $\|C_t\| < 1$ and $|q_t| < 1$. We observe this inequality in three exhaustive cases:

- (i) If $\gamma_t = 1$, then $\|C_t\| \geq 1$ and the term is 0.
- (ii) If $c_t^\Delta = 0$ and $\gamma_t \neq 1$, then the term is $(1 - \gamma_t)(1 + \gamma_t)(1 - |q_t|)$, which is strictly positive if $|q_t| < 1$ and is 0 if $|q_t| = 1$.
- (iii) If $c_t^\Delta \neq 0$ and $\gamma_t \neq 1$, then $c_t^\Delta \leq \|C_t\| = \gamma_t$, so that the term is

$$(1 - \gamma_t)[(1 + \gamma_t) - c_t^\Delta] \geq (1 - \gamma_t)1 > 0.$$

The lemma follows. □

A.3 Proof of Lemma 2

Following Lemma 1, we may assume without loss that $\|\hat{C}\| \geq 1$. Let us define our candidate $\hat{C}$. Define

$$\begin{aligned}\hat{c}_t^\Delta &:= \|C_t\| \text{ (the Euclidean norm),} \\ \hat{c}_t^\Sigma &:= 0, \\ \widehat{\Delta B}_t &:= \int_0^t \left(\frac{c^\Delta}{\|C\|} d\Delta B + \frac{c^\Sigma}{\|C\|} d\Sigma B \right) = \int_0^t \frac{1}{\hat{c}^\Delta} dq, \\ \widehat{\Sigma B}_t &:= \int_0^t \left(\frac{c^\Sigma}{\|C\|} d\Delta B + \frac{-c^\Delta}{\|C\|} d\Sigma B \right), \\ \hat{C} &:= \langle \Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P}, \hat{B}, \hat{C}, q \rangle.\end{aligned}$$

From Itô isometry, it is straightforward to see that $\frac{1}{2}\mathbb{E}_s[(B_t - B_s)(B_t - B_s)'] = (s - t)I_2$, where $I_2 \in \mathbb{R}^{2 \times 2}$ is the identity matrix. That $\frac{1}{\sqrt{2}}B$ is a standard Brownian then follows from Lévy's characterization of the same. It follows readily that $\hat{C}$ is a weak control. Moreover, that $\tau_{\hat{C}} = \tau_C$ implies

$$\begin{aligned}J(\hat{C}) - J(C) &= \mathbb{E} \int_0^{\tau_C} [(\mathbb{1}_{\hat{c}_t^\Delta=0}|q_t| + \hat{c}_t^\Delta) - (\mathbb{1}_{c_t^\Delta=0}|q_t| + c_t^\Delta)] dt \\ &= \mathbb{E} \int_0^{\tau_C} [\|C_t\| - (\mathbb{1}_{c_t^\Delta=0}|q_t| + c_t^\Delta)] dt.\end{aligned}$$

To see the value ranking, observe that the integrand has

$$\|C_t\| - (\mathbb{1}_{c_t^\Delta=0}|q_t| + c_t^\Delta) \geq \min\{\|C_t\| - c_t^\Delta, \|C_t\| - |q_t|\},$$

which is always nonnegative, and is strictly positive if $c^\Sigma \neq 0$ and $|q_t| < 1$. □

A.4 Proof of Lemma 3

The diffusion process $(q, \Delta B)$ has zero drift and volatility process $(c^\Delta, 1)$. Applying Dynkin's formula to the function $(q, \Delta B) \mapsto q_\tau \Delta B_\tau$ therefore yields $\frac{1}{2} \mathbb{E}[q_\tau \Delta B_\tau] = \mathbb{E} \int_0^\tau c_t^\Delta dt$. Moreover, Doob's optional stopping theorem tells us $\mathbb{E}[q_\tau] = q_0$. Therefore,

$$\begin{aligned} J(\mathcal{C}) &= \frac{1}{2} q_0 \Delta X_0 + \mathbb{E} \left[\int_0^\tau (\mathbb{1}_{c_t^\Delta=0} |q_t| + c_t^\Delta) dt + T - \tau \right] \\ &= 0 + T - \mathbb{E}\tau + \frac{1}{2} q_0 \Delta X_0 + \mathbb{E} \int_0^\tau c_t^\Delta dt \\ &= T - \mathbb{E}\tau + \frac{1}{2} \mathbb{E}[q_\tau] \Delta X_0 + \frac{1}{2} \mathbb{E}[q_\tau \Delta B_\tau] \\ &= T - \mathbb{E}\tau + \frac{1}{2} \mathbb{E}[q_\tau (\Delta X_0 + \Delta B_\tau)] \\ &\leq T - \mathbb{E}\tau + \frac{1}{2} \mathbb{E}|\Delta X_0 + \Delta B_\tau|, \end{aligned}$$

where the inequality is strict unless $q_\tau (\Delta X_0 + \Delta B_\tau) = |\Delta X_0 + \Delta B_\tau|$ almost surely. $\square$

A.5 Proof of Lemma 4

The arguments supporting Lemma 4 concern features of a particular optimal stopping problem.

DEFINITION 2. Given a Brownian base $\mathcal{B} = \langle \Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P}, B \rangle$ and a horizon $T \in [0, \infty]$, a $(\mathcal{B}, T)$ *stopping time* is a $[0, T]$ -valued $\{\mathcal{F}_t\}_{t \geq 0}$ stopping time.

Say a $(\mathcal{B}, T)$ stopping time is *optimal (given $(\mathcal{B}, T)$)* if it maximizes $\mathbb{E}[\frac{1}{2}|z + \Delta B_\tau| - \tau]$ over all $(\mathcal{B}, T)$ stopping times τ .

We start by proving two technical claims. The first result is that a reflected Brownian motion grows slowly enough in expectation to enable the use of various machinery from the optimal stopping literature.

CLAIM 1. Any Brownian base $\mathcal{B}$, any $z \in \mathbb{R}$, and any $\kappa > 0$ has

$$\mathbb{E} \sup_{t \in \mathbb{R}_+} (|z + \Delta B_t| - \kappa t) < \infty.$$

PROOF. Observe that

$$\begin{aligned} \mathbb{E} \sup_{t \in \mathbb{R}_+} (|z + \Delta B_t| - \kappa t) &= \mathbb{E} \max \left\{ \sup_{t \in \mathbb{R}_+} (z + \Delta B_t - \kappa t), \sup_{t \in \mathbb{R}_+} (-z - \Delta B_t - \kappa t) \right\} \\ &\leq \mathbb{E} \sup_{t \in \mathbb{R}_+} (z + \Delta B_t - \kappa t) + \mathbb{E} \sup_{t \in \mathbb{R}_+} (-z - \Delta B_t - \kappa t), \end{aligned}$$

but the latter expectations are finite. Indeed, result IV.32 from [Borodin and Salminen \(2012\)](#) implies a Brownian motion with strictly negative drift has a global maximum that is exponentially distributed and, hence, of finite mean. $\square$

The following claim adapts standard reasoning about the structure of optimal stopping problems to our specific problem. It says that the associated optimal value function is well behaved, that an optimal stopping rule exists and can be read from the optimal value function, and that the above depend only on the law governing the state rather than the specific source of randomness driving the state.

CLAIM 2. *A continuous function $v : \mathbb{R}_+ \times \mathbb{R} \rightarrow \mathbb{R}$ exists, such that for any Brownian base $\mathcal{B}$ and any $(T, z) \in \mathbb{R}_+ \times \mathbb{R}$, the $(\mathcal{B}, T)$ stopping time*

$$\tau_{T,z,v,\mathcal{B}} := T \wedge \inf \left\{ t \in [0, T] : v(T-t, z + \Delta B_t) = \frac{1}{2} |z + \Delta B_t| \right\}$$

is optimal and generates

$$\mathbb{E} \left[\frac{1}{2} |z + \Delta B_{\tau_{T,z,v,\mathcal{B}}} - \tau_{T,z,v,\mathcal{B}} \right] = v(T, z).$$

Moreover, every optimal $(\mathcal{B}, T)$ stopping time is almost surely $\geq \tau_{T,z,v,\mathcal{B}}$.

PROOF. First, fix any Brownian base $\mathcal{B}$ and let $v_{\mathcal{B}} : \mathbb{R}_+ \times \mathbb{R} \rightarrow \mathbb{R}$ be the associated optimal value function. That is, for any $(T, z) \in \mathbb{R}_+ \times \mathbb{R}$, let $v_{\mathcal{B}}(T, z)$ be the supremum of $\mathbb{E}[\frac{1}{2}|z + \Delta B_{\tau}| - \tau]$ over all $(\mathcal{B}, T)$ stopping times τ . This function is real-valued (i.e., never takes value ∞) by Claim 1. Moreover, Proposition 4.7 from [Touzi \(2012\)](#) implies $v_{\mathcal{B}}$ is continuous.²⁶

Given Claim 1 and continuity of $v_{\mathcal{B}}$, Corollary 2.9 from [Peskir and Shiryaev \(2006\)](#) implies $\tau_{T,z,v_{\mathcal{B}},\mathcal{B}}$ is optimal, thereby generating $\mathbb{E}[\frac{1}{2}|z + \Delta B_{\tau_{T,z,v_{\mathcal{B}},\mathcal{B}}} - \tau_{T,z,v_{\mathcal{B}},\mathcal{B}}] = v_{\mathcal{B}}(T, z)$. Moreover, Theorem 2.4 from [Peskir and Shiryaev \(2006\)](#) implies that any other optimal $(\mathcal{B}, T)$ stopping time is almost surely $\geq \tau_{T,z,v_{\mathcal{B}},\mathcal{B}}$.

However, now given any $(T, z) \in \mathbb{R}_+ \times \mathbb{R}$, consider any other Brownian base $\hat{\mathcal{B}}$. That $\tau_{T,z,v_{\mathcal{B}},\hat{\mathcal{B}}}$ is a $(\hat{\mathcal{B}}, T)$ stopping time implies

$$\begin{aligned} v_{\hat{\mathcal{B}}}(T, z) &\geq \mathbb{E} \left[\frac{1}{2} |z + \widehat{\Delta B}_{\tau_{T,z,v_{\mathcal{B}},\hat{\mathcal{B}}}} - \tau_{T,z,v_{\mathcal{B}},\hat{\mathcal{B}}} \right] \\ &= \mathbb{E} \left[\frac{1}{2} |z + \Delta B_{\tau_{T,z,v_{\mathcal{B}},\mathcal{B}}} - \tau_{T,z,v_{\mathcal{B}},\mathcal{B}} \right] \\ &= v_{\mathcal{B}}(T, z), \end{aligned}$$

where the first equality holds because $\mathcal{B}$ and $\hat{\mathcal{B}}$ have identical laws.

Because both $\mathcal{B}$ and $\hat{\mathcal{B}}$ were arbitrary, it follows that $v_{\mathcal{B}}$ is the same for every Brownian base $\mathcal{B}$. □

With the above two claims in place, we now proceed to prove the lemma.

²⁶That proposition refers to optimal stopping problems in which the payoff does not include a flow cost. However, the linear flow cost can be incorporated into that model by using a two-dimensional state, the second dimension tracking accrued flow cost.

PROOF OF LEMMA 4. Let $v : \mathbb{R}_+ \times \mathbb{R} \rightarrow \mathbb{R}$ be as delivered by Lemma 2, and define the set $G := \{(T, z) \in \mathbb{R}_+ \times \mathbb{R} : v(T, z) > \frac{1}{2}|z|\}$, which is relatively open in $\mathbb{R}_+ \times \mathbb{R}$ because v is continuous. For each $T \in \mathbb{R}_+$, let $G_T := \{z \in \mathbb{R} : (T, z) \in G\}$, which is open because G is. Let us make some easy starting observations about this family of sets. First, clearly, $G_0 = \emptyset$. Next, the set G_T is weakly increasing (with respect to set containment) in $T \in \mathbb{R}_+$. Indeed, v is nondecreasing in its first argument, because, for any Brownian base $\mathcal{B}$ and pair of times $t, T \in \mathbb{R}_+$ with $t \leq T$, every $(\mathcal{B}, t)$ stopping time is a $(\mathcal{B}, T)$ stopping time too. Finally, each G_T is symmetric about 0. Indeed, v is even in its second argument because, for any $(T, z) \in \mathbb{R}_+ \times \mathbb{R}$ and Brownian base $\mathcal{B} = \langle \Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P}, B \rangle$, any $(\mathcal{B}, T)$ stopping time τ is also a $(\langle \Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P}, -B \rangle, T)$ stopping time, and $\mathbb{E}[\frac{1}{2}|(-z) + (-\Delta B)_\tau| - \tau] = \mathbb{E}[\frac{1}{2}|z + \Delta B_\tau| - \tau]$.

Now we observe that every $T \in (0, \infty)$ has $G_T \ni 0$. Indeed, because T is always a $(\mathcal{B}, T)$ stopping time for any Brownian base $\mathcal{B}$, we have

$$v(T, 0) - |0| \geq \frac{1}{2}\mathbb{E}(\Delta B_T) - T = \sqrt{\frac{T}{\pi}} - T,$$

which is strictly positive for $T < \frac{1}{\pi}$. Therefore, $0 \in G_T$ for every $T \in (0, \frac{1}{\pi})$, which implies (given monotonicity of $T \mapsto G_T$) that $0 \in G_T$ for every $T \in (0, \infty)$.

Next let us see that $\bigcup_{T \in \mathbb{R}_+} G_T$ is a bounded set. To see this, we consider the relaxation of our optimal stopping problem without a deadline and apply a previously obtained solution to that time-stationary problem. Specifically, fix a Brownian base $\mathcal{B} = \langle \Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P}, B \rangle$ and let $v^* : \mathbb{R} \rightarrow \mathbb{R}$ take any $z \in \mathbb{R}$ to the supremum of $\mathbb{E}[\frac{1}{2}|z + \Delta B_\tau| - \tau]$ over all finite-mean $(\mathcal{B}, \infty)$ stopping times τ . Clearly, $v^* \geq v(T, \cdot)$ for every $T \in \mathbb{R}_+$, and so $G \subseteq \mathbb{R}_+ \times G^*$, where $G^* := \{z \in \mathbb{R} : v^*(z) > \frac{1}{2}|z|\}$. However, Theorem 16.1 from [Peskir and Shiryaev \(2006\)](#) explicitly computes the continuation region for this problem (G^* in our notation) as the set $(-\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}})$.

Finally, let us observe that G_T is convex for every $T \in (0, \infty)$. Because $G_T \ni 0$ and $\mathbb{R} \setminus G_T \supseteq (-\infty, -\frac{1}{\sqrt{2}}] \cup [\frac{1}{\sqrt{2}}, \infty)$, the property would follow if we knew both $\mathbb{R}_+ \setminus G_T$ and $\mathbb{R}_- \setminus G_T$ were convex. However, because $\frac{1}{2}|\cdot|$ is affine on $\mathbb{R}_+$ and on $\mathbb{R}_-$, the property would, in fact, follow if we knew v were (weakly) convex in its second argument. Let us now establish that fact. For any Brownian base $\mathcal{B}$, time $T \in \mathbb{R}_+$, weight $\theta \in [0, 1]$, and states $z_0, z_1 \in \mathbb{R}$, each $(\mathcal{B}, T)$ stopping time τ has

$$\begin{aligned} & \mathbb{E}\left[\frac{1}{2}|(1-\theta)z_0 + \theta z_1 + \Delta B_\tau| - \tau\right] \\ & \leq \mathbb{E}\left[\frac{1}{2}(1-\theta)|z_0 + \Delta B_\tau| + \theta|z_1 + \Delta B_\tau| - \tau\right] \\ & = (1-\theta)\mathbb{E}\left[\frac{1}{2}|z_0 + \Delta B_\tau| - \tau\right] + \theta\mathbb{E}\left[\frac{1}{2}|z_1 + \Delta B_\tau| - \tau\right] \\ & \leq (1-\theta)v(T, z_0) + \theta v(T, z_1). \end{aligned}$$

Taking the supremum over all such τ then implies $v(T, (1-\theta)z_0 + \theta z_1) \leq (1-\theta)v(T, z_0) + \theta v(T, z_1)$, as desired.

We are now ready to define $\bar{z} : \mathbb{R}_+ \rightarrow \mathbb{R}_+$. First, let $\bar{z}_0 := 0$. Then, for each $T \in (0, \infty)$, we have established that G_T is a convex open neighborhood of 0 that is symmetric about 0. That is, $G_T = (-\bar{z}_T, \bar{z}_T)$, where $\bar{z}_T := \sup G_T > 0$. Then the open set $G = \{(T, z) \in \mathbb{R}_+ \times \mathbb{R} : z < |\bar{z}_T|\}$. Moreover, our above arguments establish that $\bar{z}_T > 0$ for $T > 0$ (because $0 \in G_T$), that $\bar{z}$ is nondecreasing (because $T \mapsto G_T$ is weakly increasing with respect to set containment), and that $\bar{z}$ is bounded (because $G_T \subseteq G^* = (-\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}})$ for every $T \in \mathbb{R}_+$). The only remaining property of $\bar{z}$ to show is continuity.

Assume for a contradiction that $\bar{z}$ is discontinuous at some $T \in \mathbb{R}_+$. Because $\bar{z}$ is nondecreasing, both $\lim_{t \searrow T} \bar{z}_t$ and, if $T > 0$, $\lim_{t \nearrow T} \bar{z}_t$ exist; interpret the latter limit as $\bar{z}_0 = 0$ in the case that $T = 0$. Then let $z := \frac{1}{2} \lim_{t \searrow T} \bar{z}_t + \frac{1}{2} \lim_{t \nearrow T} \bar{z}_t$ and $\epsilon := \frac{1}{4} \lim_{t \nearrow T} \bar{z}_t - \frac{1}{4} \lim_{t \searrow T} \bar{z}_t$. So $0 < \epsilon < z$ and $\bar{z}_t$ is below $z - \epsilon$ (resp. above $z + \epsilon$) for any $t \in \mathbb{R}_+$ with $t < T$ (resp. $t > T$). Fixing a Brownian base $\mathcal{B}$, let $\tau := \inf\{t \geq 0 : |\Delta B_t| \geq \epsilon\}$. Now let v be as delivered by Claim 2 and let $\bar{v} := \max v([T, T+1] \times \{z - \epsilon, z + \epsilon\}) \in \mathbb{R}$. Then any $s \in (0, 1]$ has

$$\begin{aligned}
 2 \left[v(T+s, z) - \frac{1}{2} |z| \right] &= 2 \mathbb{E} [v(T+s - s \wedge \tau, z + \Delta B_{s \wedge \tau}) - s \wedge \tau] - z \\
 &= \mathbb{E} \{ \mathbb{1}_{\tau \geq s} [z + \Delta B_s - 2s] \} + 2 \mathbb{E} \{ \mathbb{1}_{\tau < s} [v(T+s - \tau, z + \Delta B_\tau) - \tau] \} - z \\
 &\leq \mathbb{E} \{ \mathbb{1}_{\tau \geq s} [z + \Delta B_s - 2s] \} + \mathbb{P} \{ \tau < s \} (2\bar{v}) - z \\
 &= \mathbb{P} \{ \tau < s \} (2s + 2\bar{v} - z) - 2s + \mathbb{E} [\Delta B_s - \mathbb{1}_{\tau < s} \Delta B_s] \\
 &= \mathbb{P} \{ \tau < s \} (2s + 2\bar{v} - z) - 2s + 0 - \mathbb{E} \{ \mathbb{1}_{\tau < s} \mathbb{E} [\Delta B_s | \mathcal{F}_\tau] \} \\
 &= \mathbb{P} \{ \tau < s \} (2s + 2\bar{v} - z) - 2s - \mathbb{E} [\mathbb{1}_{\tau < s} \Delta B_\tau] \\
 &\leq \mathbb{P} \{ \tau < s \} [2s + 2\bar{v} - \epsilon] - 2s.
 \end{aligned}$$

Observe now that $\tau < s$ if and only if the absolute value of the Wiener process $W := \frac{1}{\sqrt{2}} \Delta B$ exceeds $\frac{\epsilon}{\sqrt{2}}$ at some time in $[0, s]$. However, the probability of this event is no more than twice the probability that $|W_s| > \frac{\epsilon}{\sqrt{2}}$,²⁷ which is $2\Phi(\frac{-\epsilon}{\sqrt{2}s})$ because φ is even and $W_s \sim \mathcal{N}(0, \sqrt{s}^2)$. Therefore,

$$\begin{aligned}
 v(T+s, z) - \frac{1}{2} |z| &\leq \frac{1}{2} \mathbb{P} \{ \tau < s \} [2s + 2\bar{v} - \epsilon] - s \\
 &\leq 2\Phi\left(\frac{-\epsilon}{\sqrt{2}s}\right) [2s + 2\bar{v} - \epsilon] - s,
 \end{aligned}$$

but l'Hôpital's rule tells us

$$\lim_{s \rightarrow 0} \frac{\Phi\left(\frac{-\epsilon}{\sqrt{2}s}\right)}{s} = \lim_{L \rightarrow \infty} \frac{\Phi\left(\frac{-\epsilon}{\sqrt{2}}L\right)}{L^{-2}} = \frac{\epsilon}{2\sqrt{2}} \lim_{L \rightarrow \infty} \varphi\left(\frac{-\epsilon}{\sqrt{2}}L\right) L^3 = \frac{\epsilon}{4\sqrt{\pi}} \lim_{L \rightarrow \infty} e^{-\frac{\epsilon^2}{4}L^2} L^3 = 0.$$

²⁷Indeed, letting $\tilde{\tau}$ be the first time $|W|$ takes value $\frac{\epsilon}{\sqrt{2}}$, the probability that $|W_s| > \frac{\epsilon}{\sqrt{2}}$ is at least the probability that $\tilde{\tau} < s$, and $W_{\tilde{\tau}}$ lies between 0 and W_s , which is equal to half the probability of $\tilde{\tau} < s$.

Therefore, $v(T + s, z) < \frac{1}{2}|z|$ for sufficiently small $s > 0$, in contradiction to the definition of v .

Finally, we turn to establishing the uniqueness property of the optimal stopping time. Fix any Brownian base $\mathcal{B}$, any $(T, z) \in \mathbb{R}_+ \times \mathbb{R}$, and any $(\mathcal{B}, T)$ stopping time τ with $\mathbb{E}[\frac{1}{2}|z + \Delta B_\tau| - \tau]$. Letting $\tau^* := \tau_{T, z, \mathcal{B}}^*$, Claim 2 establishes that $\tau \geq \tau^*$ almost surely. Assume now, for a contradiction, that τ is not almost surely equal to τ^* . Let us observe that some $(\mathcal{B}, T)$ stopping time $\tilde{\tau} \leq \tau$ exists such that, with positive probability, $\tau > \tilde{\tau}$ and $|z + \Delta B_{\tilde{\tau}}| > \bar{z}_{T-\tilde{\tau}}$.²⁸ Now define the alternative $(\mathcal{B}, T)$ stopping times

$$\tau' := \begin{cases} \tau & : |z + \Delta B_{\tilde{\tau}}| \leq \bar{z}_{T-\tilde{\tau}}, \\ \tilde{\tau} & : |z + \Delta B_{\tilde{\tau}}| > \bar{z}_{T-\tilde{\tau}} \end{cases}$$

and $\bar{\tau} := \tau \wedge \inf\{t \in [\tau', T] : |z + \Delta B_t| \leq |\bar{z}_{T-t}|\}$. Optimality of τ implies

$$\begin{aligned} 0 &\geq \mathbb{E}\left[\frac{1}{2}|z + \Delta B_{\tau'}| - \tau'\right] - \mathbb{E}\left[\frac{1}{2}|z + \Delta B_\tau| - \tau\right] \\ &\geq \mathbb{E}\left[\frac{1}{2}|z + \Delta B_{\tau'}| - \tau'\right] - \mathbb{E}\{\mathbb{E}[v(T - \bar{\tau}, z + \Delta B_{\bar{\tau}}) - \bar{\tau} | \mathcal{F}_{\bar{\tau}}]\} \\ &= \mathbb{E}(\bar{\tau} - \tau') + \frac{1}{2}\mathbb{E}[|z + \Delta B_{\tau'}| - |z + \Delta B_{\bar{\tau}}|] \\ &= \mathbb{E}(\bar{\tau} - \tau') + \frac{1}{2}\mathbb{E}\{\mathbb{1}_{\tau > \tilde{\tau}, z + \Delta B_{\tilde{\tau}} > \bar{z}_{T-\tilde{\tau}}} \mathbb{E}[\Delta B_{\tilde{\tau}} - \Delta B_{\bar{\tau}} | \mathcal{F}_{\tilde{\tau}}]\} \\ &\quad + \frac{1}{2}\mathbb{E}\{\mathbb{1}_{\tau > \tilde{\tau}, z + \Delta B_{\tilde{\tau}} < -\bar{z}_{T-\tilde{\tau}}} \mathbb{E}[\Delta B_{\tilde{\tau}} - \Delta B_{\bar{\tau}} | \mathcal{F}_{\tilde{\tau}}]\} \\ &= \mathbb{E}(\bar{\tau} - \tau') + 0 + 0 \\ &= \mathbb{E}[\mathbb{1}_{\tau > \tilde{\tau}, |z + \Delta B_{\tilde{\tau}}| > |\bar{z}_{T-\tilde{\tau}}|} (\bar{\tau} - \tilde{\tau})] \\ &> 0, \end{aligned}$$

a contradiction. This establishes the unique optimality of τ^* (up to almost sure equality) and, hence, the lemma. $\square$

REFERENCES

- Aghion, Philippe and Matthew O. Jackson (2016), “Inducing leaders to take risky decisions: Dismissal, tenure, and term limits.” *American Economic Journal: Microeconomics*, 8 (3), 1–38. [1185]
- Benkert, Jean-Michel and Igor Letina (2020), “Designing dynamic research contests.” *American Economic Journal: Microeconomics*, 12 (4), 270–289. [1185]
- Birkinshaw, Julian (2001), “Strategies for managing internal competition.” *California Management Review*, 44 (1). [1182]

²⁸For instance, one can use $\tilde{\tau} := \tau \wedge (\frac{1}{n} + \tau^*)$ for sufficiently large $n \in \mathbb{N}$.

- Bolton, Patrick and Christopher Harris (1999), “Strategic experimentation.” *Econometrica*, 67 (2), 349–374. [1184]
- Bonatti, Alessandro and Johannes Hörner (2011), “Collaborating.” *American Economic Review*, 101 (2), 632–663. [1184]
- Bonatti, Alessandro and Heikki Rantakari (2016), “The politics of compromise.” *American Economic Review*, 106 (2), 229–259. [1185]
- Borodin, Andrei N. and Paavo Salminen (2012), *Handbook of Brownian Motion-Facts and Formulae*. Birkhäuser. [1205]
- Callander, Steven and Bård Harstad (2015), “Experimentation in federal systems.” *The Quarterly Journal of Economics*, 130 (2), 951–1002. [1185]
- Cyert, Richard M. and James G. March (1963), “A behavioral theory of the firm.” *Englewood Cliffs, NJ*, 2 (4), 169–187. [1184]
- Deb, Rahul, Mallesh M. Pai, and Maher Said (2018), “Evaluating strategic forecasters.” *American Economic Review*, 108 (10), 3057–3103. [1185]
- DeMarzo, Peter M. and Yuliy Sannikov (2006), “Optimal security design and dynamic capital structure in a continuous-time agency model.” *The Journal of Finance*, 61 (6), 2681–2724. [1185]
- Durandard, Théo (2023), “Designing dynamic research contests.” [1185]
- Farrell, Joseph and Timothy Simcoe (2012), “Choosing the rules for consensus standardization.” *The RAND Journal of Economics*, 43 (2), 235–252. [1185]
- Fudenberg, Drew, Philipp Strack, and Tomasz Strzalecki (2018), “Speed, accuracy, and the optimal timing of choices.” *American Economic Review*, 108 (12), 3651–3684. [1195]
- Gibbons, Robert (2020), “March-ing toward organizational economics.” *Industrial and Corporate Change*, 29 (1), 89–94. [1184, 1187]
- Guo, Yingni and Johannes Hörner (2020), “Dynamic allocation without money.” Northwestern University and Yale University. [1185]
- Halac, Marina, Navin Kartik, and Qingmin Liu (2017), “Contests for experimentation.” *Journal of Political Economy*, 125 (5), 1523–1569. [1184]
- Hirsch, Alexander V. and Kenneth W. Shotts (2015), “Competitive policy development.” *American Economic Review*, 105 (4), 1646–1664. [1185]
- Holmström, Bengt (1982), “Moral hazard in teams.” *The Bell Journal of Economics*, 324–340. [1184]
- Ikeda, Nobuyuki and Shinzo Watanabe (1977), “A comparison theorem for solutions of stochastic differential equations and its applications.” *Osaka Journal of Mathematics*, 14 (3), 619–633. [1189, 1202]
- Karatzas, Ioannis and Steven E. Shreve (1998), “Brownian motion.” In *Brownian Motion and Stochastic Calculus*, 47–127, Springer. [1190, 1203]

- Keller, Godfrey, Sven Rady, and Martin Cripps (2005), “Strategic experimentation with exponential bandits.” *Econometrica*, 73 (1), 39–68. [1184]
- Legros, Patrick and Hitoshi Matsushima (1991), “Efficiency in partnerships.” *Journal of Economic Theory*, 55 (2), 296–322. [1184]
- March, James G. (1962), “The business firm as a political coalition.” *The Journal of politics*, 24 (4), 662–678. [1184, 1187]
- Marino, Anthony M. and Jan Zabojnik (2004), “Internal competition for corporate resources and incentives in teams.” *RAND Journal of Economics*, 710–727. [1184]
- McClellan, Andrew (2022), “Experimentation and approval mechanisms.” *Econometrica*, 90 (5), 2215–2247. [1185]
- Mookherjee, Dilip (1984), “Optimal incentive schemes with many agents.” *The Review of Economic Studies*, 51 (3), 433–446. [1184]
- Moscarini, Giuseppe and Lones Smith (2007), “Optimal dynamic contests.” *Science (June)*, 1–18. [1185]
- Peskir, Goran and Albert Shiryaev (2006), *Optimal Stopping and Free-Boundary Problems*. Springer. [1206, 1207]
- Peters, Thomas J. and Robert H. Jr. Waterman (2003) *In Search of Excellence: Lessons From America's Best-Run Companies*, HarperCollins. [1182, 1183]
- Ryvkin, Dmitry (2022), “To fight or to give up? Dynamic contests with a deadline.” *Management Science*, 68 (11), 8144–8165. [1185]
- Sannikov, Yuliy (2008), “A continuous-time version of the principal-agent problem.” *The Review of Economic Studies*, 75 (3), 957–984. [1184, 1185, 1194]
- Stokey, Nancy L. (2016), “Wait-and-see: Investment options under policy uncertainty.” *Review of Economic Dynamics*, 21, 246–265. [1184]
- Touzi, Nizar (2012), *Optimal Stochastic Control, Stochastic Target Problems, and Backward SDE*, Vol. 29. Springer Science & Business Media. [1206]
- Yamada, Tosheo (1973), “On a comparison theorem for solutions of stochastic differential equations and its applications.” *J. Math. Kyoto Univ.*, 13 (3), 497–512. [1202]

Co-editor Bruno Strulovici handled this manuscript.

Manuscript received 5 June, 2023; final version accepted 26 February, 2025; available online 17 March, 2025.