

CURTAILING COMPETITION

Joyee Deb^{*†}
Yale University

Aditya Kuvalekar
University of Essex

Elliot Lipnowski
Columbia University

March 2021

ABSTRACT

We study project development and selection by an organization whose members prefer different projects. The organization faces a trade-off between fostering collaboration among its members and responding efficiently to projects' evolution. Choosing the most ex-post profitable project undermines members' motive to collaborate, causing ex-ante inefficiency. We show a time-varying threshold rule is uniquely optimal, with the organization choosing the first project to achieve a sufficient lead. The optimum thus entails an early phase of uninterrupted competition, always followed by a permanent collaborative regime. We develop a new martingale-based solution approach, using weak solutions to construct the optimal contract.

Keywords: project selection, internal competition, team production, collaboration, mechanism design without transfers

JEL Classification: D73, D86, D72, D82

1. Introduction

Many successful and innovative companies have been known to create internal competition between teams, assigning multiple research teams to solve the same

^{*}E-mail: joyee.deb@yale.edu, a.kuvalekar@essex.ac.uk, e.lipnowski@columbia.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, Barry Nalebuff, Harry Pei, Ben Polak, João Ramos, Debraj Ray, Tomasz Sadzik, Rani Spiegler, Balázs Szentés, 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, “Fostering Collaboration.”

technological problem or engaging distinct business units to develop competing product prototypes. This approach is especially compelling when the firm’s best way forward is not a priori clear. At the same time, sowing such internal competition rather than collaboration between teams yields potentially wasteful duplication of effort and inefficient use of resources. In their book, *In Search of Excellence: Lessons from America’s Best-Run Companies*, [Peters and Waterman Jr. \(2003\)](#) write, “Internal competition... permeates the excellent companies. It entails high costs of duplication... overlapping divisions, multiple development projects, [and] lost development dollars... Yet the benefits, though less measurable, are manifold in terms of commitment, innovation, and a focus on the revenue line.” This tension between competition and collaboration is our focus.

Consider the example of the IT 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 platform. Tech giant IBM similarly fosters competition between teams for would-be product ideas, encouraging different teams to try competing approaches to the same problem ([Peters and Waterman Jr., 2003](#)). Internal competition of this kind gives an organization flexibility at a time when what path the future will take, or which approach is the most promising, may not be clear ex ante. But this benefit of flexibility must be balanced with the costs of duplication, or the efficiency loss of wasting productive effort on the “wrong” approach.

The trade-off between the adaptive benefits and direct efficiency loss of internal competition is not unique to the firm setting. Consider the problem that party elites, across the world, face in various forms of political primaries. A literature in political science has studied the costs and benefits of candidate selection via intra-party political competition. For instance, [Adams and Merrill III \(2008\)](#) cite an important benefit: the “information-revealing advantage of holding a primary—leading to a high-quality nominee.”¹ At the same time, inducing competition via primaries entails costs: Allowing more than one candidate means inefficient nega-

¹[Adams and Merrill III \(2008\)](#) write, “[In] many plausible scenarios the strategic advantage arising from the primary electorate’s ability to select a high-quality nominee—i.e., one whose campaigning skills prove attractive to voters (such as Bill Clinton and Ronald Reagan)—outweighs the strategic disadvantage that the primary pulls the party’s nominee away from the center of the general electorate.” Relatedly, [Carey and Polga-Hecimovich \(2006\)](#) report that primaries are useful to parties in identifying candidates with higher valence, something that is hard to know ex-ante. They write, “Primaries may simply be more effective than elite-driven search processes in identifying candidates with broad popular appeal ... Carlos Menem’s emergence ... in Argentina is an example.”

tive campaigning that can harm the performance of the ultimate nominee (Carey and Polga-Hecimovich, 2008), and precious resources are spent on candidates who are ultimately unsuccessful (Adams and Merrill III, 2008).

Motivated by these applications, we study a finite-horizon game in continuous time, in which a principal interacts with two agents until a fixed deadline T . Each agent has his own project. The principal evaluates the two projects as they are developed, and must pick one of them when the deadline arrives. At every instant, an agent must decide how to allocate a unit of effort between working on his own project and providing assistance (collaborating) on the other agent’s project. Each project’s evolution is governed by a drift which is increasing in the total effort expended on that project by the two agents, and exogenous Brownian shocks. We assume that effort is costless. The vector of projects’ current state of development is publicly observed by both agents and the principal in real time, even though effort choices are not observed. 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. Our goal is to characterize the principal-optimal selection rule, assuming the principal can commit to any history-dependent project choice.²

Intuitively, to maximize her payoff, the principal would like to make the best possible project choice ex ante, foster collaboration among agents, and tailor the project choice as uncertainty resolves about which project will yield a higher payoff.³ Beginning with two simple benchmarks is useful. First, we consider the principal’s first-best policy, ignoring agency problems. Of course, the principal would then optimally defer judgment until the deadline, and pick the project with the higher final state. Moreover, at any instant, she would like both agents to allocate all of their energy to the project that is currently ahead (i.e., has a higher state at that time), to maximize the likelihood that this energy is productively useful. The first-best policy captures the intuition that the principal wants to foster collaboration, while constantly adjusting the project choice as uncertainty

²Our interest is in organizations facing real distortions that cannot be directly contracted away, leading us to a model without monetary transfers. This modeling choice reflects the classic (Gibbons, 2020; March, 1962) view 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.” Of course, understanding the nature of collaboration in the presence of transfers would be an important complementary exercise.

³Exogenous project shocks imply that an agent’s effort is solely instrumental: Effort improves project development. We abstract from another potentially important role for effort, namely, to generate more information about a given project.

resolves to ensure the agents collaborate on the “right” project. But it is not incentive compatible for the agents. Consider another natural 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. A unique equilibrium between the agents ensues, with no collaboration: Agents would rather compete, by each focusing all effort on his own project. This implies an inefficient use of the agents’ efforts.

The natural question, then, is whether it is at all possible for the principal to curtail competition, and foster some efficient collaboration. Giving the principal commitment power can help, by allowing the principal to limit herself in responding to projects’ shocks. For example, suppose that 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 perform better than 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. In fact, there is a large space of history-dependent incentive-compatible decision rules that can help the principal improve upon the no-commitment outcome, by fostering some collaboration. The main result of our paper characterizes the optimal incentive-compatible project selection rule. Our analysis yields four economic principles that guide the characterization.

1. *Foster collaboration on the “current favorite.”* Writing down the principal’s objective formally highlights that, beyond choosing an a priori promising project, the principal can increase her profit via two channels: (i) adapting the project choice as uncertainty resolves, by having the agents compete, and importantly, (ii) eliminating the efficiency loss of competition by having both agents fully collaborate on the “current favorite,” that is, the project currently more likely to be chosen by the principal (given the history and the selection rule). Determining this current favorite—which may or may not coincide with the project that has performed better to date—is a key part of characterizing the optimal selection rule.
2. *Condition only on relative performance:* We show that the principal can ignore aggregate shocks (i.e., those that the two projects both face) in her decision-making. This principle was also reflected in the first-best solution. We prove that it remains optimal even in the face of agency. The intuition is that aggregate shocks, in addition to being allocatively irrelevant given the principal’s linear objective, are completely uninformative of agent behavior.

Given such a restriction, the only way the principal resolves her eventual decision is by revising the project choice in direct response to the *difference* between the two projects' development.

3. *First compete, then switch to permanent collaboration:* Next, we show the principal optimally considers two-phase policies that involve an initial phase of agents engaging in pure competition until a stopping time, followed by both agents collaborating thereafter on the competition's winner. The intuition is that the principal's expected project choice does not change during any temporary phase of collaboration, and so she might as well backload the collaboration, better targeting effort toward the eventual choice.⁴ So the principal's problem entails an optimal stopping choice: when agents should switch from competition to collaboration.
4. *Stop competing if one project takes a large enough lead, lowering the lead threshold over time:* Finally, we find the optimal stopping rule takes an intuitive form: The principal has agents permanently switch to collaboration on the leading project when it first takes a large enough lead. Moreover, this minimal lead required to induce the principal to choose a project decreases with time. Intuitively, for a project to be chosen before the deadline, its lead over the other project must be large enough for the principal to willingly forego the potential gains from being able to adjust her project choice in the future. This option value decreases as the deadline approaches.

By deriving these guiding principles, we show that the unique principal-optimal selection rule has the following very simple form: The principal commits to a time-dependent, decreasing lead threshold $\{\hat{z}_t\}_{t \in [0, T]}$, that decreases to zero as the deadline approaches, such that a project is chosen at the first time t at which its state exceeds that of the other by at least \hat{z}_t . Equilibrium behavior therefore also has a simple pattern: In an initial competitive regime, agents allocate all effort toward their own project. This phase is followed by a collaborative phase, with agents permanently collaborating until the deadline, on the project that is first to achieve the threshold lead. Moreover, this threshold lead decreases over time, capturing the diminishing benefits of option value from competition. In particular, regardless of the time horizon, a non-degenerate phase of collaboration always exists.

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.

Notice that agents are indifferent between competing and collaborating in the collaboration phase. This feature may make the reader wonder whether an agent could ever be willing to collaborate if he thought there were any chance at all that his project would be chosen in the future. In other words, is the irreversible switch to collaboration merely an artefact of agent indifference during collaboration? Perhaps surprisingly, the answer is negative. The two-phase rule with a permanent switch to collaboration is a feature of the second-best *solution*. Indeed, there exist many incentive-compatible decision rules that involve agents switching back and forth between competition and collaboration. Agents are willing to collaborate even if they assign positive probability to being selected in the future based on the performance of their respective projects. But in deriving the third guiding principle described above, we show that such rules turn out to be suboptimal.⁵

Returning to our motivating applications, we see this two-phase optimal contract is consistent with what we observe. Indeed, competing teams are eventually brought together, once and for all, toward a common approach, once enough relative uncertainty has been resolved. At Telstar, top-level executives finally chose in favor of EX, and the two teams subsequently collaborated on EX in order to build a common platform for the future. At IBM, teams are allowed to work on disparate approaches until, at some point, the firm conducts performance “shootouts” among the competing groups to pick one (Peters and Waterman Jr., 2003). Likewise, consistent with our model’s prediction of a collapsing threshold, American presidential primaries consistently produce a presumptive nominee before concluding. In the speech announcing she was suspending her campaign for the 2008 US Democratic nomination, Hillary Rodham Clinton publicly urged her supporters that it was time for her and them “to take [their] energy, [their] passion, [their] strength, and do all [they] can to help elect Barack Obama.”⁶ Following an intense competition, Clinton joined the Obama effort, telling her supporters, “[Work] as hard for Barack Obama as you have for me.”

Let us briefly highlight our approach to deriving the optimal selection rule. The principal needs to choose a project at the deadline as a function of the entire history of projects’ evolution. The space of such history-contingent selection rules is unwieldy. Our continuous-time model enables us to somewhat simplify the principal’s problem. In particular, every selection rule induces a natural martingale—expected eventual project choice given the history up to the current time. This martingale helps us in two ways. First, we can recast the principal’s problem as

⁵See Section 3.3 for a more detailed discussion on this point.

⁶See <https://www.nytimes.com/2008/06/07/us/politics/07text-clinton.html>.

an easier stochastic control problem, where (even though the true design problem is a static decision based on a rich history) we can interpret the principal as dynamically deciding how to respond to contemporaneous shocks.⁷ Second, because the agents' interests are opposed, their incentives are now straightforward: They compete if the expected project choice responds to the projects' current relative performance, and are willing to collaborate otherwise. The principal's control problem amounts to trading off tailoring the project choice based on project performance with efficiently harnessing gains from collaboration on the eventually-chosen project. However, given our finite-horizon setting (which would render the standard Hamilton-Jacobi-Bellman equation a partial differential equation), even the existence of an optimal selection rule is not immediate. Moreover, even if existence were guaranteed, the non-stationary character of the principal's optimization problem would stand in the way of deriving the intuitive qualitative features that we discuss above. To circumvent this issue, we employ methods based on (i) a focus on more permissive *weak solutions* of stochastic differential equations and (ii) the known equivalence between Itô integrals and time-changed Brownian motions, to construct the type of controls described above to establish our main result.

A key consequence of our characterization of the optimal policy is that the initial phase of competition is always temporary: Regardless of the horizon, the phase of collaboration is reached with probability 1. In particular, the unique optimal contract implies arbitrarily large ex-post inefficiencies can occur on path. On the one hand, fostering collaboration increases the value of the principal's chosen project because both agents work on it. On the other hand, having a long phase of collaboration makes large mistakes prevalent. In other words, for any given margin, at the deadline, with positive probability, the project that is not chosen ends up ahead by more than that margin. But, intuitively, the probability of such mistakes is lower when the duration of collaboration is longer.

Finally, consistent with our motivating applications, we could consider a contracting environment in which the principal has the option of irreversibly canceling a project at any time, after which both agents must work on the remaining project. Our project selection rule and consequent agent behavior can be implemented in equilibrium in such a setting, and remains optimal. This alternate formulation highlights that the commitment our principal requires is quite weak: Rather than

⁷This same recursive approach is, in principle, available even in discrete time. However, our continuous formulation makes the principal's space of choices of how to respond to contemporaneous shocks considerably simpler, just as it does in, say, [Sannikov \(2008\)](#).

committing to an arbitrary history-contingent project selection rule, the principal need only commit to irreversible project choices.⁸

The remainder of the paper is organized as follows. First, we survey some of the most related literature below. In section 2, we present the formal model. We start with presenting two simple benchmarks in section 3: the first-best policy (ignoring agent incentives) and the case with no commitment by the principal. In section 4, we more concretely characterize agent incentive compatibility, conceptually clarifying the trade-off our principal faces. Section 5 contains the main results of the paper. In section 6, we discuss some consequences of our main result. Proofs, other than that of our main theorem, are in the Appendix.

1.1. Related Literature

At a high level, our paper is indebted to the perspective of organizations as political coalitions, as outlined by March (1962) and Cyert and March (1963). That work highlights that individuals within organizations have goals that are often distinct from the goals of the organization, and the role of the executive is that of a political broker who cannot solve such problems by simple payments.⁹ We focus on one such conflict between individuals and the organization—that different members would like the organization to support their own pet initiatives, which may undermine collaboration.

More specifically, our principal’s problem can be interpreted as a multi-agent experimentation problem. The literature on experimentation in teams is large, some prominent examples being Bolton and Harris (1999), Keller et al. (2005), and Bonatti and Hörner (2011). Given the compete-then-collaborate form of our optimal contract, a notable point of reference in this literature is Halac et al. (2017), who study optimal contest design to promote innovation. Although many other modeling differences exist between these papers and our own, the most important is that agents in these other papers face an incentive to free ride on each others’ costly experimentation.¹⁰ In our model, with each agent choosing only how to allocate a fixed effort budget between two projects, and agents’ interests

⁸Notice that this alternate implementation is viable only because of the two-phase structure, with backloaded collaboration, that we show is optimal. An incentive-compatible rule that involved agents switching back and forth between competing and collaborating could not have been implemented with such little commitment power.

⁹See Gibbons (2020) for a detailed survey.

¹⁰Free-riding in teams is an extensively studied topic outside the experimentation framework as well. For instance, somewhat related to our work, Marino and Zabožnik (2004) show how internal competition can be beneficial in addressing the free-rider problem.

being directly opposed, free-riding is not a concern. The central trade-off in our work is between retaining option value via competition and harnessing gains from collaboration. Another key point of contrast to the experimentation literature is that, in our model, the flow of information—the variance reduction concerning each project’s final state—is exogenous to agents’ choices. Hence, while our principal faces an exploration-exploitation tradeoff as players do in that literature, its source is distinct. In experimentation models, a decision maker trades off the information generated from exploration against the myopic value of exploitation that it crowds out. In the present model, information arrives exogenously, but when the principal uses it to inform future choices, she distorts current effort.

The closest work to our own is that of [Bonatti and Rantakari \(2016\)](#). There, each of two agents first chooses what type of project to develop and how hard to work in developing it over time, after which they negotiate over the adoption choice. In their framework, projects differ in the payoffs generated for both players, whose interests are partially aligned. The focus of that work, therefore, is on studying the nature of projects agents choose to develop, and the resulting negotiations that ensue. A key lesson is that the mechanism by which projects are selected can feed into the development stage when agents wish to distort the organization’s decision, a feature shared by the static work of [Hirsch and Shotts \(2015\)](#) and the two-period political model of [Callander and Harstad \(2015\)](#).¹¹ This lesson sets the stage for our design problem.

Finally, our paper can potentially contribute methodologically to the growing literature on dynamic mechanism design without transfers (e.g., [Aghion and Jackson, 2016](#); [Deb et al., 2018](#); [Guo and Hörner, 2020](#)). Within that literature, future allocative decisions may optimally be distorted in response to current shocks, either to provide present incentives (as in [Meyer, 1992](#)) or in response to learning or other payoff-relevant shocks (as in [Meyer, 1991](#)). A recent contribution featuring both such sources of dynamics is [McClellan \(2020\)](#), whose principal faces a hypothesis-testing problem subject to interim participation constraints for a privately informed agent. Like our work, [McClellan \(2020\)](#) employs tools from the literature on dynamic contracting in continuous time (e.g., [DeMarzo and Sannikov, 2006](#); [Sannikov, 2008](#)), and he also must study his principal’s sequential problem directly to circumvent analyzing a partial differential equation. We are hopeful that the specific techniques employed in our paper—appealing directly to martingale methods rather than Hamilton-Jacobi-Bellman equations to reduce the

¹¹See also [Farrell and Simcoe \(2012\)](#), who study related distortions in standards adoption across firms that produce complementary products.

principal’s control problem, and passing between weak and strong solutions—will enable the use of continuous-time contracting methods more broadly in dynamic mechanism design without transfers.

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 is the “owner” of one 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. Let $a_t^i \in [0, 1]$ denote the fraction of effort that agent i allocates to his own project at time t . He 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. Effort is costless and contributes to projects’ development continuously over time. Formally, 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 $\langle \Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P} \rangle$ satisfying the usual conditions; 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 $\langle \Omega, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P} \rangle$; and project i has exogenous initial state $X_0^i \in \mathbb{R}$. The vector of project states is publicly observed by both agents and the principal, and effort-allocation choices are not observed. It is convenient to define the following:

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

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 $\langle \Omega, \mathcal{F}, \mathbb{P} \rangle$ for a payoff of $\mathbb{E}[X_T^y]$; 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 $X_0^1 + X_0^{-1} = 0$.¹² Therefore, after normalization, project X^i follows

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

and the principal's payoff is

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

Each agent wants her project to be chosen; that is, agent i gets payoff $\mathbb{E}[iy]$. 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 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 $\ell_t := \operatorname{argmax}_{i \in I} X_t^i$.

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. It turns out that the first-best solution has the agents collaborating at every instant, while in the equilibrium with no principal commitment, it is not possible to get any collaboration off the ground. These two results lead us to ask whether it is possible for a principal with some commitment power to curtail competition in order to get better equilibrium outcomes.

3.1. First-best solution: Ignoring agent incentives

Toward defining the principal's first-best problem formally, let \mathcal{A} denote the set of $[0, 1]^2$ -valued progressively measurable processes on $\{\mathcal{F}_t\}_t$ and let \mathcal{Y} denote the set of $\{-1, 1\}$ -valued random variables on \mathcal{F} . We want to solve the following planner

¹²Counting time in different units, we may assume without loss that $\sigma = \mu$; then, positively rescaling the state (and so principal payoffs), we may take $\sigma = \mu = 1$; and finally, as adding a constant to both projects' initial states or both projects' drifts simply adds a constant to the principal's payoff (not altering any incentives or the quality of information about agents' behavior), we may further assume $\beta = -1$ and $X_0^1 + X_0^{-1} = 0$. When we state our main theorem, we clarify exactly how these parameters alter the form of the optimal selection rule.

problem:

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

The proposition below shows that 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. In the interim, we show 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:*

- *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 principal chooses project $y^{FB} = \ell_T$, the leader as of time T .*

The intuition for this result is straightforward. Because the principal will optimally choose the ex-post best project, her objective can be rewritten as $\frac{1}{2}\mathbb{E}|\Delta X_T|$, an increasing transformation of $(\Delta X)^2$. But then the given control increases the drift of $(\Delta X)^2$ more than any other control does, at any given level of $(\Delta X)^2$. Hence, a classic comparison theorem from the theory of stochastic differential equations 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 (just 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 being

collaborative) weakly increases agent i 's payoff in every state, strictly so with positive probability if he was not already almost surely making the latter choice at almost every time.

PROPOSITION 2: *If the principal cannot commit, then:*

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

3.3. Can commitment be useful?

Absent commitment power, the principal is perfectly responsive, and the result is no collaboration. 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 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.

Notice 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, suppose that the principal commits to the following more elaborate selection rule that gives an early leader an advantage, though the advantage is not unassailable. If a project i is the first to take on a lead of L then the principal chooses this early leader i with probability p at the deadline. However, if subsequently, project $-i$ catches up so that agent i 's lead is reduced to 0, and $-i$ remains ahead at the deadline, then

project i is chosen only with probability $2p - 1$. However, if project i is again ahead at the deadline, then 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.

We can show that this selection rule is also incentive-compatible, and can dominate the outcomes under both the no-commitment and the unassailable lead contracts.¹³

The key takeaway from these examples is that commitment power can indeed help the principal trade off some of the efficiency benefits from curtailed competition against the benefits of better ex-post choice from increased competition. Naturally, the space of such history-dependent contracts is large, and the substance of our main result is to identify the uniquely optimal one.

4. Agent Incentives and the Principal's Problem

Recall, the principal in our model 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. Toward better understanding the principal's problem, we now express agent incentive compatibility 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 interim expected project choice. By the martingale representation theorem (Karatzas and Shreve, 1998, Theorem 3.4.15), a progressively measurable \mathbb{R}^2 -valued process on $\{C_t = (c_t^\Delta, c_t^\Sigma)\}_t$ on filtration $\{\mathcal{F}_t\}_{t \geq 0}$ exists whose time- t quadratic variation has finite expectation for every $t \geq 0$, and such that

$$q_t = q_0 + \int_0^t \left[c_t^\Delta (d\Delta X_t - 2\widetilde{\Delta a}_t dt) + c_t^\Sigma d\Sigma X_t \right], \quad (1)$$

where $\widetilde{\Delta a}_t$ is the equilibrium-anticipated Δa_t , and the law of motion $d\Delta X_t$ is influenced by the actually-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

¹³Observe, the early-advantage lead contract demonstrates that an agent can be willing to collaborate even when he knows that, with sufficient improvement, his project can be chosen in the future. We discuss this possibility again in Section 5.4

to relative and aggregate shocks of the two projects. It is immediate from the expression above that for agent incentive compatibility, we must have

$$\Delta a_t = 0 \text{ whenever } c_t^\Delta \neq 0. \quad (2)$$

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 as follows:

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

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 the correct ex-ante project choice, the principal's two levers to increase profit are (i) adapting the project choice by having the agents compete and work on their respective project, and (ii) eliminating the efficiency loss of competition by having agents collaborate.

Imposing the agent incentive-compatibility constraint (2), we can derive an upper bound for the principal's objective:

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

Recall that q_t is the interim expected project choice. So, we can see from the above expression, that to get efficient collaboration, the principal should have the agents fully collaborate on the "current favorite," that is, the project currently more likely to be chosen by the principal. Note that the current favorite may not be the project that is currently ahead. Determining what the current favorite is at any time is an essential part of the characterization of the optimal contract.

We can now write down a relaxed version of the principal's problem:¹⁴

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

$$\text{where } q_t \text{ is given by (1), or equivalently, } dq_t = c_t^\Delta d\Delta B_t + c_t^\Sigma d\Sigma B_t. \quad (3)$$

5. The Optimal Selection Rule

The following theorem, our main result, describes the form of the uniquely optimal project selection rule.

THEOREM 1: *A bounded, continuous, nondecreasing 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 is optimal:*

- *Each agent works on his own project before $\tau^* := \inf\{t \in [0, T] : |\Delta X_t| \geq \bar{z}_{T-t}\}$,¹⁵*
- *The principal chooses project $y^* = \ell_{\tau^*}$, the leader as of time τ^* ;¹⁶*
- *Both agents work on project y^* from time τ^* onward.*

Moreover, this optimum is unique: Any other optimal incentive-compatible selection rule almost surely has the same chosen project and the same agent choices at almost every time.

The remainder of the section is dedicated to proving the theorem.

One reasonable 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 that such a solution is in fact the principal's optimal value function. However, this direct approach has two limitations. First, given the finite horizon, the HJB equation would be a PDE, and so establishing the existence of a smooth solution to it is not straightforward. Moreover, we are interested in understanding qualitative properties of the optimal selection rule, and these properties would be hard to show without an explicit

¹⁴Because we eventually show this optimal value is attainable with an incentive-compatible contract, the given augmentation of the principal's objective will turn out to be payoff irrelevant.

¹⁵In the version of the model described before we normalized several parameters, we let $\tilde{z}_t := \frac{\sigma^2}{\mu} \bar{z} \frac{\mu^2}{\sigma^2 t}$ and take τ^* to be the first time $t \in [0, T]$ at which $|\Delta X_t| \geq \tilde{z}_{T-t}$.

¹⁶In the zero-probability event that $\tau^* = T$ and $X_T^1 = X_T^{-1}$, the principal may choose arbitrarily.

characterization of such a solution to the PDE, even if its existence were guaranteed. So, we adopt a different route. Because the argument is somewhat involved and not typical of the optimal contracting literature, we describe the main steps of the proof in section 5.1 below: That subsection does not contain any formal arguments, but rather focuses on explaining our approach. The formal results are presented thereafter.

5.1. Approach to characterizing the optimal selection rule

In this subsection, we informally describe our proof approach.

Our first technical step is to consider relaxations of the principal’s problem that allow for *weak solutions*.¹⁷ 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 that, now, we 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). In a typical discrete-time model, such a relaxation would be irrelevant, but in the present setting, it turns out to be a useful tool for the analyst. Given this broader definition of a control, we then proceed to show restricting attention to controls that have some economically intuitive features is without loss of optimality.

In Lemma 1, we show that two properties are without loss of optimality: First, the principal ignores aggregate shocks. Given our earlier observation that c^Σ does not affect the principal’s objective function or agent incentives, and that aggregate shocks are not an informative signal of agents’ choices, it is intuitive that the principal should set $c^\Sigma = 0$. Second, the principal sets $c^\Delta > 0$ and has the agents compete until either the principal settles her choice of project or the deadline arrives. Notice from (3) that q_t is unchanged during a phase of collaboration, and so conjecturing that the principal cannot be worse off if she backloads collaboration to a time when it may be better targeted is reasonable.

In Lemma 2, we show we can reduce the principal’s (relaxed) 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

¹⁷Although permitting weak solutions is the most nonstandard sense in which we relax the principal’s problem, it is not the only one. Additionally, we consider only certain necessary conditions for agent incentives, and use an objective function that is generally only an upper bound on the principal’s objective. As is typical, when combining these arguments yields an incentive-compatible contract at which the augmented objective coincides with the true objective, it then follows that such a contract is optimal.

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 from then on.

Finally, in Lemma 3, we show the stopping rule is a decreasing threshold:¹⁸ 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.

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, in the sense that the projects' Brownian shock process provides the same payoff for this incentive-compatible policy as the principal could have attained under the weak solution.

5.2. Mathematical preliminaries

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

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

- (i) $\langle \Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P} \rangle$ is a filtered probability space satisfying the usual conditions;
- (ii) $B = (\Delta B, \Sigma B) = \{B_t\}_{t \geq 0}$ is a \mathbb{R}^2 -valued stochastic process on $\{\mathcal{F}_t\}_{t \geq 0}$ such that $\frac{1}{\sqrt{2}}B$ is a standard Brownian motion;
- (iii) $C = (c^\Delta, c^\Sigma) = \{C_t\}_{t \geq 0}$ 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) $q = \{q_t\}_{t \geq 0}$ is a $[-1, 1]$ -valued martingale on $\{\mathcal{F}_t\}_{t \geq 0}$;
- (v) $q_t = q_0 + \int_0^t C \cdot dB$ 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.

Notice that if a principal must optimally choose a control defined as 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

¹⁸Studying costly sequential sampling problems that a single decision-maker faces, [Fudenberg et al. \(2018\)](#) show decreasing threshold rules can arise naturally even without a deadline.

Brownian base $\langle \Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P}, B \rangle$ as given. But it is convenient to consider this relaxation of the principal's problem.

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

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

Intuitively, given a 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 $\tau_{\mathcal{C}}$.

Notice expression (O) implies $J(\mathcal{C})$ is an upper bound on the payoff of the principal. In what follows, we consider the optimal control problem with $J(\mathcal{C})$ as the objective.

5.3. Only relative performance matters

We first 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 this property remains optimal even respecting agent incentives. Furthermore, having the agents compete until either the principal makes a project choice or time runs out is without loss of optimality.

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

The interested reader can refer to the Appendix for the proof, but we summarize the logic here. The proof proceeds in two steps.

We first establish a quantitative claim: For the principal to resolve uncertainty somewhat quickly is without loss of optimality. Specifically, if she is deciding slowly enough about which project to choose that its flow benefits are surely smaller than those from collaboration on a chosen project, the principal may as well speed up this resolution of uncertainty and defer any saved time toward endgame collaboration. Formally, this claim amounts to showing that restricting attention to controls such that $\|\hat{C}\| \geq 1$ is without loss of optimality. The proof is constructive, modifying a control without this property to a superior one with

this property.

Specifically, the fractal property of Brownian motion allows us to construct a superior control, by replacing the underlying Brownian motion with a law-equivalent time change of the same, and 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 “speeding up 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 fixed an expected project choice q_0 , the principal’s payoff in (O) is a sum of the total net value of competition ($\int_0^T c_t^\Delta dt$) and the total accrued value of collaboration ($\int_0^T \mathbb{1}_{c_t^\Delta=0} |q_t|$). 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. So the cost of lost flow payoff as a result of speeding up is always less than the benefit of the extra collaboration time.

In the next step, we show that restricting attention to 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. The proof is considerably less involved than that of the previous step. 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 does without loss by the first step), these costs are smaller than the gains to more effective competition.

5.4. First compete, then switch to collaboration

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 2: *For any control \mathcal{C} such that $c^\Delta \geq 1$ and $c^\Sigma = 0$, the stopping time $\tau :=$*

$\tau_{\mathcal{C}}$ has $J(\mathcal{C}) \leq T + \mathbb{E} \left[\frac{1}{2} |\Delta X_0 + \Delta B_\tau| - \tau \right]$. Moreover, $J(\mathcal{C}) < T + \mathbb{E} \left[\frac{1}{2} |\Delta X_0 + \Delta B_\tau| - \tau \right]$ 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 in the Appendix.

At this point, the reader might suspect the two-phase structure—in which the principal makes an irreversible project choice and then fosters collaboration—is closely linked to the fact that agents are indifferent in any collaboration phase. In particular, it would be reasonable to conjecture that collaboration can never be induced without abandoning one of the projects once and for all. The early-lead advantage contract described in Section 3.3 shows that such a conjecture is false. Under the early-lead advantage contract, in equilibrium, agents start out competing, then collaborate on the project with an early lead, and then can switch back to competing. In other words, it is possible to foster collaboration even if an agent remains uncertain about which project will be adopted, and indeed even if some of this uncertainty will be resolved in response to projects' future performance.

5.5. Stop competing when the lead is large, lowering the standard continuously

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 3: *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*

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

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 $T > 0$.*

The proof is in the Appendix, but we sketch the argument here. We have an optimal stopping problem, where a decision-maker observes a driftless Brownian motion at a constant flow cost, and can stop at any time before a deadline, where

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.

stopping yields a payoff equal to its absolute value. The finite deadline makes the problem non-stationary, 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 uniquely optimal policy is to stop as soon as the optimal and stopping values coincide. So, we analyze the (continuous) optimal value function, taking as arguments the time remaining and the current state of the Brownian motion, and show 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 it contains zero when near enough to the deadline, and hence (given monotonicity) contains zero at every time. It is symmetric about zero because the objective and law of motion are. Finally, it is an interval around zero because the value function is convex, whereas the terminal value is affine on either side of zero.

5.6. Characterizing the optimal selection rule

The qualitative insights in the preceding subsections were derived for weak solutions to the principal's problem. However, they 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 3, let Π^* be the principal value generated by the behavior named in the theorem. Observe 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 control \mathcal{C} such that the $J(\mathcal{C}) \geq \Pi$.

Now, let us apply each of our three building blocks. Lemma 1 delivers some 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 2 then tells

us the stopping time $\hat{\tau} := \tau_{\hat{c}}$ has $T + \mathbb{E} \left[\frac{1}{2} |\Delta X_0 + \widehat{\Delta B}_{\hat{\tau}}| - \hat{\tau} \right] \geq \Pi$, 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 3 tells us τ^* (as defined in the statement of the theorem) has $T + \mathbb{E} \left[\frac{1}{2} |\Delta X_0 + \Delta B_{\tau^*}| - \tau^* \right] \geq \Pi$, strictly so unless $\hat{\tau}$ is almost surely equal to $\tau_{T,z,\hat{c}}^*$.

The above arguments directly deliver the theorem. First, they show the principal’s optimal value is $\Pi^* = T + \mathbb{E} \left[\frac{1}{2} |\Delta X_0 + \Delta B_{\tau^*}| - \tau^* \right]$, 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. Because the switching threshold decreases as the deadline approaches, the probability of such mistakes is lower for projects that have experienced longer collaboration.

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.

6.2. Cancellation of projects before the deadline

In our setting, the principal chooses an optimal stopping time at which she makes a permanent project 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. To allow the principal to irreversibly cancel one of the projects is perhaps more consistent with our two motivating applications of product development within organizations and candidate selection via political primaries. Observe that this richer contracting environment is equivalent to the one we have studied. To see the equivalence, first note the principal could not be worse off in the richer environment, because she could always abstain from canceling projects before the deadline. Conversely, the principal can always simulate project cancellation through a selection rule by (just as in our Theorem 1) deciding on a project in advance, and having the agents collaborate on the chosen project (which is incentive compatible even though working on the other project is an available choice). Thus, the two are equivalent.

Although equivalent for optimal-contracting purposes, the model with irreversible termination enables a simple implementation of our optimal selection rule (and consequent agent behavior) in Theorem 1: 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 canceled. Our optimum would remain optimal in this alternate setting, and would indeed also be sequentially rational. Put differently, in an alternate model in which the principal can commit to irreversibly cancel a project, additional commitment power would not benefit her, because our optimal selection rule could be implemented in equilibrium.

6.3. Agent indifference in the collaboration phase

In the optimal contract, during the collaboration phase, agents are indifferent between competing and collaborating. This sort of indifference is common in many mechanism-design problems, but many such environments have the feature that constructing a similar contract with strict incentives is possible.²¹

In our setting, no such contract is available. With no transfers, no observability,

²¹For 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.

and directly opposed interests between the agents, the principal has very few instruments at her disposal. Indeed, the best that the principal can do in any strict equilibrium is simply the no-commitment solution (the principal chooses the leading project at the deadline and agents always compete).²² However, if the principal could additionally provide monetary incentives up to a fixed budget of $\epsilon > 0$ (assuming for simplicity that agents have separable preferences over money and project choice), then it would be easy to implement this optimum in a strict equilibrium. 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 $\frac{\epsilon}{2}$ 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 support the chosen project in the collaborative phase, and it does not distort incentives in the competitive phase because an agent's monetary prize of $\frac{\epsilon}{2}$ will be earned with probability $\frac{1}{2}$ conditional on any outcome of the initial competition. Moreover, one can show the principal's optimal value converges to ours as $\epsilon \rightarrow 0$.

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 some collaboration in equilibrium is still possible and optimal despite the paucity of powerful incentivizing instruments.

References

- J. Adams and S. Merrill III. Candidate and party strategies in two-stage elections beginning with a primary. *American Journal of Political Science*, 52(2):344–359, 2008.
- P. Aghion and M. O. Jackson. Inducing leaders to take risky decisions: dismissal, tenure, and term limits. *American Economic Journal: Microeconomics*, 8(3): 1–38, 2016.

²²Given 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.

- J. Birkinshaw. Strategies for managing internal competition. *California Management Review*, 44(1), Fall 2001.
- P. Bolton and C. Harris. Strategic experimentation. *Econometrica*, 67(2):349–374, 1999.
- A. Bonatti and J. Hörner. Collaborating. *American Economic Review*, 101(2): 632–63, 2011.
- A. Bonatti and H. Rantakari. The politics of compromise. *American Economic Review*, 106(2):229–59, 2016.
- A. N. Borodin and P. Salminen. *Handbook of Brownian motion-facts and formulae*. Birkhäuser, 2012.
- S. Callander and B. Harstad. Experimentation in federal systems. *The Quarterly Journal of Economics*, 130(2):951–1002, 2015.
- J. M. Carey and J. Polga-Hecimovich. Primary elections and candidate strength in latin america. *The Journal of Politics*, 68(3):530–543, 2006.
- J. M. Carey and J. Polga-Hecimovich. The primary elections ‘bonus’ in latin america. *Designing democratic government*, pages 227–247, 2008.
- R. M. Cyert and J. G. March. A behavioral theory of the firm. *Englewood Cliffs, NJ*, 2(4):169–187, 1963.
- R. Deb, M. M. Pai, and M. Said. Evaluating strategic forecasters. *American Economic Review*, 108(10):3057–3103, 2018.
- P. M. DeMarzo and Y. Sannikov. Optimal security design and dynamic capital structure in a continuous-time agency model. *The Journal of Finance*, 61(6): 2681–2724, 2006.
- J. Farrell and T. Simcoe. Choosing the rules for consensus standardization. *The RAND Journal of Economics*, 43(2):235–252, 2012.
- D. Fudenberg, P. Strack, and T. Strzalecki. Speed, accuracy, and the optimal timing of choices. *American Economic Review*, 108(12):3651–84, 2018.
- R. Gibbons. March-ing toward organizational economics. *Industrial and Corporate Change*, 29(1):89–94, 2020.

- Y. Guo and J. Hörner. Dynamic allocation without money. *Northwestern University and Yale University*, 2020.
- M. Halac, N. Kartik, and Q. Liu. Contests for experimentation. *Journal of Political Economy*, 125(5):1523–1569, 2017.
- A. V. Hirsch and K. W. Shotts. Competitive policy development. *American Economic Review*, 105(4):1646–64, 2015.
- N. Ikeda and S. Watanabe. A comparison theorem for solutions of stochastic differential equations and its applications. *Osaka Journal of Mathematics*, 14(3):619–633, 1977.
- I. Karatzas and S. E. Shreve. Brownian motion. In *Brownian Motion and Stochastic Calculus*, pages 47–127. Springer, 1998.
- G. Keller, S. Rady, and M. Cripps. Strategic experimentation with exponential bandits. *Econometrica*, 73(1):39–68, 2005.
- J. G. March. The business firm as a political coalition. *The Journal of politics*, 24(4):662–678, 1962.
- A. M. Marino and J. Zabojnik. Internal competition for corporate resources and incentives in teams. *RAND Journal of Economics*, pages 710–727, 2004.
- A. McClellan. Experimentation and approval mechanisms. *Unpublished working paper*, 2020.
- M. A. Meyer. Learning from coarse information: Biased contests and career profiles. *The Review of Economic Studies*, 58(1):15–41, 1991.
- M. A. Meyer. Biased contests and moral hazard: Implications for career profiles. *Annales d’Economie et de Statistique*, pages 165–187, 1992.
- G. Peskir and A. Shiryaev. *Optimal stopping and free-boundary problems*. Springer, 2006.
- T. J. Peters and R. H. Waterman Jr. *In Search of Excellence: Lessons from America’s Best-Run Companies*. HarperCollins, 2003.
- Y. Sannikov. A continuous-time version of the principal-agent problem. *The Review of Economic Studies*, 75(3):957–984, 2008.

N. L. Stokey. Wait-and-see: Investment options under policy uncertainty. *Review of Economic Dynamics*, 21:246–265, 2016.

T. Yamada. On a comparison theorem for solutions of stochastic differential equations and its applications. *J. Math. Kyoto Univ.*, 13(3):497–512, 1973.

7. Appendix: Omitted Proofs

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

7.1. Proof of Proposition 1

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

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

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

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

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 control has a (weakly) first-order-stochastically dominated distribution of $|\Delta X_T|$.²³ \square

7.2. Proof of Lemma 1

In order to prove Lemma 1, we first prove two claims. The first claim says that it is without loss of optimality for the principal to resolve uncertainty somewhat quickly. Our argument is a “replication” argument that couples realized paths of

²³That result shows a control $\Delta a_t = -\text{sign}(\Delta X_t)$ minimizes $|\Delta X_T|$ —in fact, minimizes each of $\{|\Delta X_t|\}_{t \in [0, T]}$ —in an FOSD sense. However, reproducing the proof nearly verbatim establishes that control $\Delta a_t = \text{sign}(\Delta X_t)$ maximizes $|\Delta X_T|$.

play from the new control to those from the old control, traversed at a different “speed.” Such an argument is only available in continuous time, due to the self-similarity properties of Brownian motion.

CLAIM 1: *For any control \mathcal{C} , a control $\hat{\mathcal{C}}$ exists whose Euclidean norm satisfies $\|\hat{\mathcal{C}}\| \geq 1$ 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 $|q_t| < 1$.*

Proof. Let $\tau := \tau_{\mathcal{C}}$, 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 zero on any time interval where it is a.e. zero.

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

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

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)

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

To see that $\hat{\mathcal{C}}$ is a 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 $\int_0^u \hat{C} \cdot d\hat{B} = \int_0^{\lambda_u} C \cdot dB$. Let us defer this property until the end of the proof, and first show the desired value ranking holds if this stochastic differential equation holds.

Taking for granted that $\hat{\mathcal{C}}$ is a control, we now proceed to show that $J(\hat{\mathcal{C}}) \geq J(\mathcal{C})$. 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

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

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

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

The value ranking will then follow if we establish that $(1 - \gamma_t) \left[(1 + \gamma_t)(1 - \mathbb{1}_{c_t^\Delta=0} |q_t|) - c_t^\Delta \right]$ 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:

1. If $\gamma_t = 1$, then $\|C_t\| \geq 1$ and the term is zero.
2. 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 zero if $|q_t| = 1$.
3. 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

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

We return now to our one unresolved detail: showing that $\int_0^u \hat{C} \cdot d\hat{B} = \int_0^{\lambda_u} C \cdot dB$, which will (because $\|\hat{C}\| \geq 1$ by construction) establish the claim. Letting M^λ be the local martingale on $\{\hat{\mathcal{F}}_u\}_{u \geq 0}$ given by $M_u^\lambda := \int_0^{\lambda_u} C \cdot dB$, it will be useful to consider the \mathbb{R}^3 -valued local martingale \vec{M} on $\{\hat{\mathcal{F}}_u\}_{u \geq 0}$ given by

$$\vec{M}_u := \begin{pmatrix} \widehat{\Delta B}_u \\ \widehat{\Sigma B}_u \\ M_u^\lambda \end{pmatrix} = \int_0^{\lambda_u} \begin{pmatrix} \gamma \, d\Delta B \\ \gamma \, d\Sigma B \\ C \cdot dB \end{pmatrix}.$$

By direct computation, and using the fact that $\frac{1}{\sqrt{2}}(\widehat{\Delta B}, \widehat{\Sigma B})$ is a standard Brownian motion, the quadratic covariation (matrix) process of \vec{M} up to time $u \geq 0$ is given by

$$\langle \vec{M} \rangle_u = 2 \begin{pmatrix} u & 0 & \int_0^{\lambda_u} \gamma c^\Delta \\ 0 & u & \int_0^{\lambda_u} \gamma c^\Sigma \\ \int_0^{\lambda_u} \gamma c^\Delta & \int_0^{\lambda_u} \gamma c^\Sigma & \int_0^{\lambda_u} \|C\|^2 \end{pmatrix}.$$

Toward further simplifying the above expression, consider any process $\xi \in \{\gamma c^\Delta, \gamma c^\Sigma, \|C\|^2\}$. Then, interpreting the expression $\frac{\xi}{\gamma^2}$ arbitrarily wherever $\xi = \gamma = 0$, observe that

$$\int_0^{\lambda_u} \xi = \int_0^{\lambda_u} \frac{\xi}{\gamma^2} \, d\zeta = \int_0^{\lambda_u} \frac{\xi \lambda_\zeta}{\gamma_{\lambda_\zeta}^2} \, d\zeta = \int_0^{\zeta_{\lambda_u}} \frac{\xi_\lambda}{\gamma_\lambda^2} = \int_0^u \frac{\xi_\lambda}{\gamma_\lambda^2}.$$

Substituting in the definition of \hat{C} , it follows that

$$\langle \vec{M} \rangle_u = 2 \int_0^u \begin{pmatrix} 1 & 0 & \hat{c}^\Delta \\ 0 & 1 & \hat{c}^\Sigma \\ \hat{c}^\Delta & \hat{c}^\Sigma & \|\hat{C}\|^2 \end{pmatrix}.$$

Now, defining the local martingales $\hat{M}, \hat{\hat{M}}$ on $\{\hat{\mathcal{F}}_u\}_{u \geq 0}$ via $\hat{M}_u := \int_0^u \hat{c}^\Delta \, d\widehat{\Delta B}$ and $\hat{\hat{M}}_u := \int_0^u \hat{c}^\Sigma \, d\widehat{\Sigma B}$, our goal is to show the process $\hat{M} + \hat{\hat{M}} - M^\lambda$ is almost surely zero. But because the process is a local martingale, it suffices to show its quadratic variation is zero. And indeed,

$$\begin{aligned} \frac{1}{2} \langle \hat{M} + \hat{\hat{M}} - M^\lambda \rangle_u &= \frac{1}{2} \langle \hat{M} \rangle_u + \frac{1}{2} \langle \hat{\hat{M}} \rangle_u + \frac{1}{2} \langle M^\lambda \rangle_u + \langle \hat{M}, \hat{\hat{M}} \rangle_u - \langle \hat{M}, M^\lambda \rangle_u - \langle \hat{\hat{M}}, M^\lambda \rangle_u \\ &= \int_0^u [(\hat{c}^\Delta)^2 + (\hat{c}^\Sigma)^2 + \|\hat{C}\|^2] + 0 - \int_0^u [\hat{c}^\Delta \, d\langle \widehat{\Delta B}, M^\lambda \rangle + \hat{c}^\Sigma \, d\langle \widehat{\Sigma B}, M^\lambda \rangle] \\ &= 2 \int_0^u \|\hat{C}\|^2 - 2 \int_0^u [\hat{c}^\Delta (\hat{c}^\Delta) + \hat{c}^\Sigma (\hat{c}^\Sigma)] \\ &= 0, \text{ as required.} \end{aligned}$$

□

Given the previous claim, the following claim shows that for the principal to respond only to the relative performance of projects, not to aggregate shocks, and moreover to respond positively to a project's relative performance until the decision is fully made or the deadline arrives, is without loss of optimality. The formal argument for the claim follows a similar (though less involved) “replication” argument to that of the previous claim. Here, rather than changing the “clock” of the underlying Brownian motion, we change the “angle” by altering which principal component of a given shock is attributable to each project.

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

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

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

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 $\hat{\mathcal{C}}$ is a control. Moreover, that $\tau_{\hat{\mathcal{C}}} = \tau_{\mathcal{C}}$ implies

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

To see the value ranking, observe that the integrand has

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

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

Finally, observe that Lemma 1 follows immediately from Claim 2. \square

7.3. Proof of Lemma 2

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,

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

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

7.4. Proof of Lemma 3

The arguments supporting Lemma 3 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 $\mathbb{E} \left[\frac{1}{2} |z + \Delta B_\tau| - \tau \right]$ over all (\mathcal{B}, T) -stopping times τ .

We start with 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 3: Any Brownian base \mathcal{B} , any $z \in \mathbb{R}$, and any $\kappa > 0$ have

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

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 one. It says 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 4: *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*

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

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 $\mathbb{E} \left[\frac{1}{2} |z + \Delta B_{\tau} | - \tau \right]$ over all (\mathcal{B}, T) -stopping times τ . This function is real-valued (i.e., never takes value ∞) by Claim 3.

Let us observe that $v_{\mathcal{B}}$ is continuous. To see this, consider any $(T, z), (\tilde{T}, \tilde{z}) \in \mathbb{R}_+ \times \mathbb{R}$. For any (\mathcal{B}, T) -stopping time τ , it is immediate that $\tau \wedge \tilde{T}$ is a (\mathcal{B}, \tilde{T}) -

stopping time. Therefore,

$$\begin{aligned}
& \mathbb{E} \left[\frac{1}{2} |z + \Delta B_\tau| - \tau \right] - v_{\mathcal{B}}(\tilde{T}, \tilde{z}) \\
& \leq \mathbb{E} \left[\frac{1}{2} |z + \Delta B_\tau| - \tau \right] - \mathbb{E} \left[\frac{1}{2} |\tilde{z} + \Delta B_{\tau \wedge \tilde{T}}| - \tau \wedge \tilde{T} \right] \\
& = \frac{1}{2} \mathbb{E} [|z + \Delta B_\tau| - |\tilde{z} + \Delta B_{\tau \wedge \tilde{T}}|] - \mathbb{E} [\tau - \tau \wedge \tilde{T}] \\
& \leq \frac{1}{2} \mathbb{E} [|z + \Delta B_\tau| - |\tilde{z} + \Delta B_{\tau \wedge \tilde{T}}|] \\
& \leq \frac{1}{2} \mathbb{E} |(z + \Delta B_\tau) - (\tilde{z} + \Delta B_{\tau \wedge \tilde{T}})| \\
& \leq \frac{1}{2} |z - \tilde{z}| + \frac{1}{2} \mathbb{E} |\Delta B_\tau - \Delta B_{\tau \wedge \tilde{T}}| \\
& \leq \frac{1}{2} |z - \tilde{z}| + \frac{1}{2} \mathbb{E} \left| \Delta B_{(\tau \wedge \tilde{T}) + |T - \tilde{T}|} - \Delta B_{\tau \wedge \tilde{T}} \right| \\
& \leq \frac{1}{2} |z - \tilde{z}| + \frac{1}{\sqrt{\pi}} \sqrt{|T - \tilde{T}|}.
\end{aligned}$$

Taking the supremum over all such τ then implies

$$v_{\mathcal{B}}(T, z) - v_{\mathcal{B}}(\tilde{T}, \tilde{z}) \leq \frac{1}{2} |z - \tilde{z}| + \frac{1}{\sqrt{\pi}} \sqrt{|T - \tilde{T}|}.$$

Because this inequality holds for all such pairs, the function $v_{\mathcal{B}}$ is continuous.

Given Claim 3 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 $\mathbb{E} \left[\frac{1}{2} |z + \Delta B_{\tau_{T,z,v_{\mathcal{B}},\mathcal{B}}}| - \tau_{T,z,v_{\mathcal{B}},\mathcal{B}} \right] = v_{\mathcal{B}}(T, z)$. Moreover, Theorem 2.4 from [Peskir and Shiryaev \(2006\)](#) implies that any other optimal (\mathcal{B}, T) -stopping time is almost surely $\geq \tau_{T,z,v_{\mathcal{B}},\mathcal{B}}$.

But 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 B and \hat{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} . \square

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

Proof of Lemma 3. Let $v : \mathbb{R}_+ \times \mathbb{R} \rightarrow \mathbb{R}$ be as delivered by Lemma 4, 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 zero. 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 $\mathbb{E} \left[\frac{1}{2} |(-z) + (-\Delta B)_\tau| - \tau \right] = \mathbb{E} \left[\frac{1}{2} |z + \Delta B_\tau| - \tau \right]$.

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 $\mathbb{E} \left[\frac{1}{2} |z + \Delta B_\tau| - \tau \right]$ over all finite-mean (\mathcal{B}, ∞) -stopping times τ . Clearly, $v^* \geq v(T, \cdot)$ for every $T \in \mathbb{R}_+$, and so $G \subseteq \mathbb{R}_+ \times G^*$, where $G^* := \{z \in \mathbb{R} : v^*(z) > \frac{1}{2}|z|\}$. But 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 $\mathbb{R} \setminus G_T \supseteq (-\infty, -\frac{1}{2}] \cup [\frac{1}{2}, \infty)$, the property would follow if we knew both $\mathbb{R}_+ \setminus G_T$ and $\mathbb{R}_- \setminus G_T$ were convex. But, 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 zero that is symmetric about zero. 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 Lemma 4, 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 Wiener process $W := \frac{1}{\sqrt{2}} \Delta B$ exceeds $\frac{\epsilon}{\sqrt{2}}$ at some time in $[0, s]$. But the probability of this event is no more than twice the probability that $|W_s| > \frac{\epsilon}{\sqrt{2}}$,²⁴ which is $2\Phi\left(\frac{-\epsilon}{\sqrt{2s}}\right)$ because φ is even and $W_s \sim \mathcal{N}(0, \sqrt{s^2})$. Therefore,

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

²⁴Indeed, 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 $\tilde{\tau} < s$.

But L'Hôpital's rule tells us

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

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 4 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 $|z + \Delta B_{\tilde{\tau}}| > \bar{z}_{T-\tilde{\tau}}$.²⁵ But then, defining 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

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

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

²⁵For instance, one can use $\tilde{\tau} := \tau \wedge (\frac{1}{n} + \tau^*)$ for sufficiently large $n \in \mathbb{N}$.