

FOSTERING COLLABORATION

Joyee Deb^{*†}
Yale University

Aditya Kuvalekar
University of Essex

Elliot Lipnowski
Columbia University

June 2023

ABSTRACT

We study a basic trade-off that organizations face: competition versus collaboration. Maintaining competing approaches provides option value when the best choice is uncertain. Collaborating on a single approach prevents inefficient use of resources. We model project development and selection by a principal interacting with two agents who prefer their respective project. Opposed interests undermine agents' incentive to collaborate, causing inefficiencies. We show a time-varying threshold rule is uniquely optimal: the principal selects the first project to achieve a sufficient lead. The optimum entails initial competition, always followed by permanent collaboration. Our proof uses new martingale methods applying weak solutions.

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

JEL Classification: D73, D86, D72, D82

*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, Stephen Morris, Barry Nalebuff, Harry Pei, Ben Polak, João Ramos, Debraj Ray, Tomasz Sadzik, Rani Spiegler, Balázs Szentes, Yu Fu Wong, and Bill Zame for helpful feedback. We are especially grateful to Mykhaylo Shkolnikov for his advice. We also thank various conference and seminar participants, especially Martin Cripps and Meg Meyer, who served as discussants for this paper in the ERC 2019 Frontiers in Design conference and the NBER Fall 2020 Organizational Economics Working Group, respectively. A previous draft of this paper was circulated with the title “Curtailing Competition.” Aditya Kuvalekar gratefully acknowledges financial support by MICIN/ AEI/10.13039/501100011033 grant (PGC2018-096159-B-I00) from the National Plan, PID2019-107259RB-I00 grant from Agencia Estatal de Investigación del Gobierno de España, and by MICIN/ AEI/10.13039/501100011033 grant (PID2021-123747NA-I00), which he received while he was at Universidad Carlos III de Madrid.

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 [and] innovation...

[Peters and Waterman Jr. \(2003\)](#)

In Search of Excellence: Lessons from America's Best-Run Companies

1. Introduction

In the face of uncertainty, organizations often create internal competition to determine the best course of action. Maintaining and developing competing approaches provides option value, when the best course of action is initially uncertain. For example, political parties use primaries to select a preferred candidate. A large literature in political science has studied the costs and benefits of candidate selection via intra-party political competition. A party has difficulty predicting a priori how public sentiment toward a candidate—particularly the valence or non-policy related appeal of a candidate—may change over time and which candidate will be ultimately more attractive to the electorate.¹ Holding primaries allows a party to adapt its final choice to changing public sentiment. [Adams and Merrill III \(2008\)](#) cite this “information-revealing advantage of holding a primary—leading to a high-quality nominee.”² At the same time, inducing competition via primaries entails costs: allowing more than one candidate means precious resources are spent on candidates who are ultimately unsuccessful ([Adams and Merrill III, 2008](#)).

Using internal competition to adapt to an uncertain environment is not unique to political organizations. Many successful companies create internal competition between teams, assigning multiple research teams to solve the same technological problem or engaging distinct business units to develop competing product prototypes. 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

¹[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.”

²[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.”

worked on its own platform, knowing the firm would ultimately adopt exactly one platform.³ When the most promising approach is not clear ex ante, sowing competition rather than collaboration between teams gives an organization flexibility to adapt its choice based on what the future will look like. But the benefit of flexibility must be balanced with the costs of duplication, or the efficiency loss of wasting productive effort on the “wrong” approach.

This tension between the adaptive benefits and the efficiency losses of competition inside organizations is at the heart of our paper. Our central question is how to design optimal selection rules that harness both the flexibility gains of competition and the efficiency gains of collaboration.

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

Intuitively, the principal strives to achieve the following three objectives: making the best possible project choice ex ante, fostering collaboration among agents (make them work on the same project), and tailoring the final choice as uncertainty resolves about which project will yield a higher payoff.

First, we show that in the principal’s first-best policy, ignoring agency problems, the principal would simply wait until the deadline and pick the project with the higher final state. Moreover, at any instant, she would like both agents to collaborate, that is, to allocate all of their effort to the project that is currently ahead (i.e., has a higher state at that time), to maximize the likelihood that this effort is

³Tech 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).

productively useful. The first-best policy captures the intuition that the principal wants to foster collaboration, while constantly adjusting the project choice as uncertainty resolves, to ensure the agents collaborate on the “right” project. But it is not incentive compatible for the agents, because they will each strictly prefer to work on their own project. Consider another benchmark, in which the principal can make no prior commitments regarding the eventual project choice. It is easy to see that, in this case too, the principal will pick the project with the higher ex-post state at the deadline. A unique equilibrium between the agents ensues, with no collaboration. Agents would rather compete, each focusing all effort on their own project. The result is an inefficient use of the agents’ efforts.

The natural question, then, is whether the principal can curtail competition and foster some efficient collaboration. Giving the principal commitment power can help. For example, suppose the principal commits to choosing the project that is first to take a lead by a specified margin. It is easy to see how such a rule can outperform the outcome under no commitment (pure competition). Moreover, it is incentive compatible, because once an agent’s project is chosen irreversibly, agents are indifferent and thus willing to collaborate. At this point, one might wonder whether the only way for the principal to induce agents to collaborate is to make an irreversible choice at some point. One may reasonably conjecture that with strictly opposing interests, if an agent thought their project could be adopted in the future after enough improvement, however small the likelihood, they would strictly prefer to work on their own project. Perhaps surprisingly, this conjecture is false. Many incentive-compatible decision rules induce agents to switch back and forth between competing and collaborating.⁴ Indeed, the space of history-contingent selection rules that can improve upon the no-commitment outcome by fostering some collaboration is rich and unwieldy, and our main result characterizes the optimal incentive-compatible selection rule.

We show the unique principal-optimal selection rule has a simple form: the principal commits to a time-dependent, decreasing lead threshold $\{\hat{z}_t\}_{t \in [0, T]}$ that decreases to zero as the deadline approaches, such that a project is chosen at the first instant t at which its state exceeds that of the other by at least \hat{z}_t . Agent equilibrium behavior therefore also has a simple pattern, namely, that in an initial competitive regime (before a project is chosen by the principal), agents allocate all effort toward their own project. This phase is followed by permanent collaboration; that is, agents collaborate on the chosen project (the first to achieve the threshold lead) until the deadline. In particular, regardless of the time horizon, a non-degenerate phase of collaboration always exists. The main force behind

⁴See section 3.3 for a more detailed discussion on this point.

this two-phase structure is the option value arising from frontloading competition: sustaining temporary collaboration before additional competition necessitates that the principal not change her expected project choice during the collaboration phase. In that case, however, she can do better by frontloading competition, saving collaboration for when it can be better targeted. The decreasing lead threshold captures the diminishing benefits of option value from competition.

This two-phase optimal contract with inevitable collaboration from some point onward is broadly consistent with what we observe in our motivating applications. American presidential primaries typically produce a presumptive nominee well before the deadline (the last primary), after which everyone supports the chosen nominee.⁵ Similarly, competing teams in organizations are eventually brought together toward a common approach once enough relative uncertainty has been resolved. At Telstar, top-level executives finally chose EX over AX, and both teams subsequently collaborated on EX to build a common platform for the future.⁶

The derivation of the optimal contract proceeds in four logical steps. We first show any contract is outperformed by one in which *the principal resolves her decision quickly*. Intuitively, if she does not adapt her current choice adequately to project shocks for some time, she is not utilizing the option value of competition. Speeding up her decision makes room for additional collaboration targeted to the eventual chosen project.⁷ Second, we show *the principal conditions only on relative performance*, because aggregate shocks are both irrelevant to the principal’s objective and uninformative about agent behavior. Third, we show it is optimal for the principal to consider two-phase policies with *initial competition until a stopping time, followed by a permanent switch to collaboration*; in particular, she makes a constrained-efficient choice with the partial information available when ending competition, and has both agents collaborate thereafter on the chosen winner. Finally, we show the principal optimally *chooses a project only when its lead over the other project is sufficiently large, lowering her standards closer to the deadline*. Intuitively, when ending competition and choosing a project, the principal foregoes the option value from adjusting to projects’ future shocks, but this

⁵In her 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” (see <https://www.nytimes.com/2008/06/07/us/politics/07text-clinton.html>). Following an intense competition, Clinton joined the Obama effort, telling her supporters, “[Work] as hard for Barack Obama as you have for me.”

⁶At IBM, teams are allowed to work on disparate approaches until, at some point, the firm conducts performance “shootouts” to pick one (Peters and Waterman Jr., 2003).

⁷A 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.

option value vanishes as the deadline approaches.

A key feature of the optimal policy is that the initial phase of competition is always temporary. Regardless of the horizon, collaboration starts strictly before the deadline, with probability 1. This feature implies arbitrarily large ex-post inefficiencies can occur on path. In other words, for any given margin, with positive probability, the project that is not chosen ends up ahead at the deadline, by more than that margin. But the probability of such mistakes vanishes as the duration of collaboration grows long.

Finally, briefly discussing our methodological contribution is worthwhile. Although the optimal selection rule and its qualitative properties are intuitive in hindsight, deriving them is not easy, because of the richness of the space of possible rules. Our continuous-time model is useful in simplifying the problem. Every selection rule induces a natural martingale—the expected eventual project choice given the current history—which allows us to recast the principal’s problem as an easier stochastic control problem, in which the principal is deciding dynamically how to respond to contemporaneous shocks, even though the true design problem is a static decision based on a rich history.⁸ However, given our finite-horizon setting, even the existence of an optimal control is not immediate. Even if existence were guaranteed, the non-stationarity of the problem makes deriving qualitative features of the optimal rule difficult. So, instead of using the standard toolkit of Hamilton-Jacobi-Bellman equations, we adopt a different approach. We consider a particular relaxation of the principal’s problem, allowing for more permissive *weak solutions* of stochastic differential equations. Weak solutions enable us to use the known equivalence between Itô integrals and time-changed Brownian motions to construct optimal controls with the economically intuitive features described above. We hope this approach of appealing directly to martingale methods to simplify the principal’s problem will be applied more broadly in economic theory.

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 \(1963\)](#), which 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 po-

⁸This same approach is, in principle, available even in discrete time. Our continuous-time formulation makes the principal’s space of choices of how to respond to contemporaneous shocks considerably simpler, just as it does in, say, [Sannikov \(2008\)](#).

litical broker who cannot solve such problems by simple payments.⁹ We focus on one such conflict—that different members would like the organization to support their own pet initiatives, which may undermine collaboration.

Our paper is related to the work on multi-agent experimentation. See, for instance, [Bolton and Harris \(1999\)](#), [Keller et al. \(2005\)](#), [Bonatti and Hörner \(2011\)](#), and [Halac et al. \(2017\)](#). The fundamental trade-offs are different though. In those papers, agents face an incentive to free-ride on each others' costly experimentation.¹⁰ Free-riding is absent from our model. Our central trade-off is between retaining option value via competition and harnessing gains from collaboration. Another important difference is that, in our model, the flow of information is exogenous and unaffected by agents' choices. Hence, even though our principal faces an exploration-exploitation trade-off as players do in that literature, its source is quite 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 our paper, information arrives exogenously, but when the principal uses it to inform future choices, she distorts the agents' current effort choice. In our first-best solution, no such trade-off exists: the principal can both exploit (induce collaboration on the leading project) and explore (adapt choice to the projects' evolution). But doing so violates agent incentives. So, our exploration-exploitation trade-off arises endogenously, exactly because of the agency problem.

The closest work to ours is that of [Bonatti and Rantakari \(2016\)](#), where each agent first chooses what type of project to develop and how hard to work on it over time, after which they negotiate over the adoption choice. Agent interests are partially aligned. The focus is on studying the nature of projects agents choose to develop and the negotiations that ensue. A key lesson is that the project-selection mechanism can feed into the development stage when agents may distort the organization's decision. This feature also arises in the static model of [Hirsch and Shotts \(2015\)](#) and the two-period model of [Callander and Harstad \(2015\)](#).¹¹ This lesson sets the stage for our design problem.

A large literature on dynamic contests focuses on moral hazard—effort provision—in contests while abstracting away from collaboration (e.g., [Benkert and Letina, 2020](#); [Moscarini and Smith, 2007](#); [Ryvkin, 2022](#)). These papers are primarily concerned, therefore, with eliciting high effort from the players. In our work, with

⁹See [Gibbons \(2020\)](#) for a detailed survey.

¹⁰Free-riding in teams is an extensively studied topic outside the experimentation framework as well. For example, see [Holmström \(1982\)](#), [Mookherjee \(1984\)](#), and [Legros and Matsushima \(1991\)](#). More thematically related to our work, [Marino and Zabochnik \(2004\)](#) show how internal competition can be beneficial in addressing the free-rider problem.

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

no margin for total effort, we can focus on incentivizing the agents to direct their efforts towards the better project improving efficiency.

Finally, our paper contributes 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](#)). Both these sources of dynamics arise in [McClellan \(2021\)](#), whose principal faces a hypothesis-testing problem subject to interim participation constraints for a privately informed agent. Like us, [McClellan \(2021\)](#) employs tools from the literature on dynamic contracting in continuous time (e.g., [DeMarzo and Sannikov, 2006](#); [Sannikov, 2008](#)) and studies his principal’s sequential problem directly to circumvent analyzing a partial differential equation. We are hopeful our techniques—appealing to martingale methods rather than HJB equations to reduce the principal’s control problem, and passing between weak and strong solutions—will be used more extensively 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 has a project with evolving state X^i . The principal must pick one of the two projects at the deadline T . At every instant, each agent must allocate a unit of effort between working on their 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 their own project at time t . Agent i allocates the remaining $(1 - a_t^i)$ of their effort to helping agent $-i$ on their 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. For convenience, define

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

and define ΔB , ΣB , and Δa analogously.

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 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 expected payoff is

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

Each agent wants her project to be chosen; that is, agent i gets payoff iy . Given any (y, a^1, a^{-1}) , we can define $q_t^i := \mathbb{E}[iy|\mathcal{F}_t]$ as agent i 's continuation value at any $t \leq T$. In what follows, we write agent incentives from the point of view of agent 1. So, we drop the superscript i and define $q_t := q_t^1 = \mathbb{E}[y|\mathcal{F}_t]$. We denote the current leader at any time $t \leq T$ by $\ell_t := \operatorname{argmax}_{i \in I} X_t^i$.

2.1. Discussion of assumptions

Let us discuss some of our modeling assumptions.

Unobservable effort choice: We assume the agent's choice of effort allocation is unobservable to the principal—or at least non-contractible. In the political application, this assumption is reasonable. Political parties are decentralized and do not closely monitor candidate effort, and much of the effort taken by politicians is behind-the-scenes work. In the case of a firm, we interpret effort as creative effort. Many modern technology companies foster innovation and develop new solutions by giving employees the flexibility to spend a fraction of their time on projects of their choice. We view effort in our model as such unmonitored time that agents can choose to use on their own or another agent's pet project.

¹²Counting time in different units, we may assume without loss that $\sigma = \mu$; then, positively rescaling the state (and thus principal payoffs), we may take $\sigma = \mu = 1$; and finally, because 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.

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

Fixed agent payoff, no transfers: We assume an agent gets no payoff from the value of the chosen project itself. This assumption is again consistent with the leading application, in which all that the candidate cares about in the primary stage is being the chosen nominee. In a firm, we interpret the agents’ preference for having their project chosen as akin to empire-building motives, with individuals seeking to increase their scope and influence within organizations.¹³ If we allowed the principal to offer agents a share of the chosen project or to make transfers, inducing collaboration would be easier. Our optimal rule highlights how, despite a severe lack of instruments, the principal can still foster some collaboration.¹⁴

Exogenous information: We assume information is exogenous. This assumption is appropriate for settings in which effort is mainly instrumental. For instance, political candidates can affect their chances of success by devoting more or less effort to campaigning, but campaigning effort allocation does not help predict how public sentiment may change in the future. We abstract from the potential information-generating role of effort. As mentioned in section 1.1, the literature on experimentation focuses on this role and models an exploitation-exploration trade-off that is technological: exploring a risky option means giving up the myopic gains from exploiting the safe option. Our modeling choice enables us to make the different point that, even if information arrives exogenously, which means a single decision-maker would face no trade-off between exploration (competition)

¹³See <https://www.investopedia.com/terms/e/empirebuilding.asp> for more on empire-building.

¹⁴We are interested in organizations facing distortions that cannot be directly contracted away. 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.”

and exploitation (collaboration), using this information to inform future choices can still entail costly distortions because of agent incentives.

Continuous time: Our choice of a finite-horizon continuous-time model allows us to simplify the optimization problem and derive qualitative features of the optimal decision rule. The reader may wonder, after seeing the simple form of the optimal rule, whether it could not have been generated in a two-period model. It turns out that solving a two-period model is not clean, and moreover does not yield some economically substantive insights that our model delivers. For instance, a two-period model cannot demonstrate the somewhat unintuitive point that the principal can get agents to switch back and forth between competing and collaborating in an incentive-compatible way, despite their directly opposed interests—or the point that inducing the agents to do so is strictly suboptimal. A two-period model also cannot deliver the insight that the optimal rule involves a bounded phase of competition followed by collaboration, with most of the time being spent collaborating when the duration is long.

3. Benchmarks

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

3.1. First-best solution: Ignoring agent incentives

Toward defining the principal’s first-best problem 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 problem:

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

The proposition below shows the first-best solution is for the principal to choose the project with the higher output at the deadline and, at every instant before the deadline, have both agents collaborate on the current leader. One part is obvious: the principal will clearly choose the better project ex post. We also show that at

any instant before the deadline, it is optimal to have the agents collaborate on the current best guess of which project will be ultimately chosen, so that the effort is productive. Formally, we observe that it is optimal to set $\Delta a_t = 1$ when $\Delta X_t > 0$ and $\Delta a_t = -1$ when $\Delta X_t < 0$.

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

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

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

The intuition for this result is straightforward. Because the principal will optimally choose the ex-post best project, her objective can be rewritten as $\frac{1}{2}\mathbb{E}|\Delta X_T|$, 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 $(\Delta X)^2$. A classic comparison theorem from the theory of stochastic differential equations (Ikeda and Watanabe, 1977) says this control yields a stochastically maximal distribution of $(\Delta X_T)^2$.

3.2. No principal commitment

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

PROPOSITION 2: *If the principal cannot commit:*

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

3.3. Can commitment be useful?

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

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

This rule maximizes the benefits of collaboration on the favored project for length of time T , completely forgoing the benefits of choosing the correct project. It is easy to see 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.

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

Early-Lead Advantage: Finally, consider the following more elaborate selection rule that gives an early leader an advantage, though the advantage is not unassailable. If a project i is the first to take on a lead of L , the principal chooses this early leader i with probability $p \geq \frac{1}{2}$ at the deadline. If project $-i$ subsequently catches up so that agent i 's lead is reduced to 0, and $-i$ remains ahead at the deadline, project i is chosen only with probability $2p - 1$. However, if project i is again ahead at the deadline, i is chosen with probability 1. Under this rule, agents start out competing. When an early leader i emerges, both agents start collaborating on the early-leader project and start competing again if and when the early lead disappears.

This selection rule is incentive compatible—in particular, agents willingly collaborate when an early leader emerges because the probability of the early

leader being chosen at the deadline, conditional on the lead disappearing, is still $\frac{1}{2}(1) + \frac{1}{2}(2p - 1) = p$. Moreover, such a rule can dominate the outcomes under both the no-commitment and the unassailable-lead contracts.

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

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

4. Agent Incentives and the Principal's Problem

Recall the principal can choose an arbitrary $\{-1, 1\}$ -valued random variable y on $\langle \Omega, \mathcal{F}, \mathbb{P} \rangle$, and her expected payoff $\mathbb{E}[X_T^y]$ depends on the resulting agent behavior. So, we start with expressing 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 the 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 that¹⁵

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

where $\widetilde{\Delta a}_t$ is the equilibrium-anticipated Δa_t , and the law of motion $d\Delta X_t$ is influenced by the chosen Δa_t . Intuitively, we can think of c_t^Δ and c_t^Σ as project sensitivities that describe how the interim expected project choice responds to relative and aggregate shocks of the two projects. It is immediate from the expression above that for agent incentive compatibility, we must have

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

¹⁵That is, $q_t = q_0 + \int_0^t \left[c_{\tilde{t}}^\Delta (d\Delta X_{\tilde{t}} - 2\widetilde{\Delta a}_{\tilde{t}} d\tilde{t}) + c_{\tilde{t}}^\Sigma d\Sigma X_{\tilde{t}} \right]$ almost surely.

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

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

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 has two levers to increase profit: (i) adapting the project choice to relative productivity shocks, which will (by (2)) induce agents to compete on their respective projects; 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—indeed, such collaboration whenever c_t^Δ is zero is the only way to attain the upper bound. Note that the current favorite may not be the project that is currently ahead, because the former is endogenous to the principal's chosen contract. Determining what the current favorite is at any time is an essential part of the characterization of the optimal contract.

Hence, we can write the principal's problem as:

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

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

5. The Optimal Selection Rule

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

THEOREM 1: *An optimal contract exists and is unique.¹⁶ 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:*

- *The principal chooses project $y^* = \ell_{\tau^*}$, the leader as of time¹⁷*

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

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

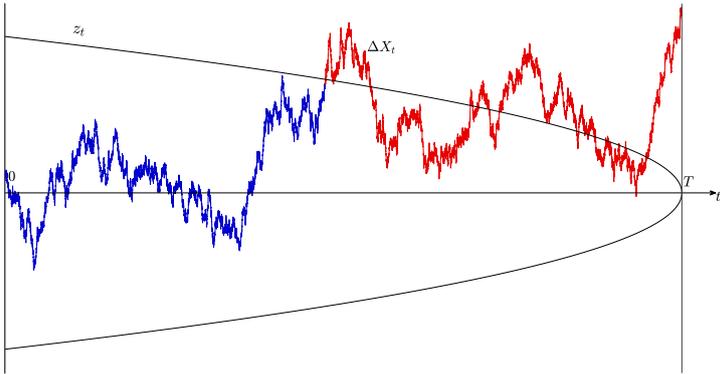


Figure 1: Project 1 is selected

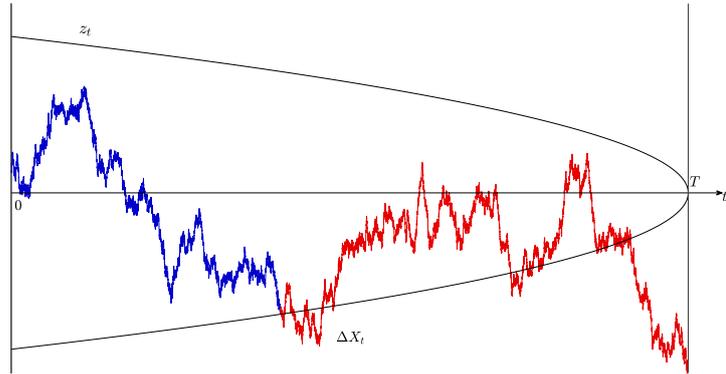


Figure 2: Project -1 is selected

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

¹⁶Any two optimal incentive-compatible selection rules almost surely have the same chosen project and the same agent choices at almost every time.

¹⁷In the zero-probability event that $\tau^* = T$ and $X_T^1 = X_T^{-1}$, the principal may choose arbitrarily.

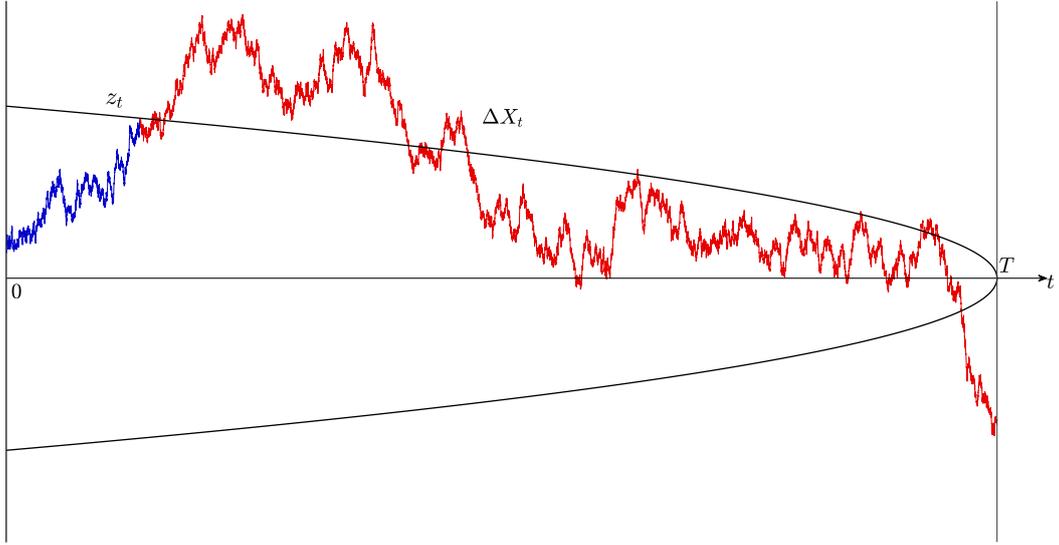


Figure 3: An ex-post inefficient choice

Recall, the theorem is stated with the normalizations $\sigma = \mu = 1, \beta = -1$ and $\Sigma X_0 = 0$. Given arbitrary σ, μ, β , and ΣX_0 , it is easy to deduce that agents compete until the first instant \tilde{t} at which one project is leading by a margin of at least

$$\tilde{z}_{\tilde{t}} := \frac{\sigma^2}{\mu} \bar{z} \left(\frac{\mu^2}{\sigma^2} \right) (T - \tilde{t}),$$

at which point that project is chosen and collaborated upon until the deadline. In particular, the initial aggregate state (ΣX_0) and baseline development rate (β) are irrelevant, for the same reasons that aggregate performance is optimally ignored. Because each quantity enters the threshold rule in two places, comparative statics with respect to the marginal product of effort (μ) and the project volatility (σ) are more delicate, depending on detailed features of the function \bar{z} . However, some features are easy: for instance, raising μ while holding the ratio μ/σ fixed raises the lead threshold at every time. Finally, extending the deadline (T) leads to a larger lead threshold at every time, whereas changing the initial project asymmetry (ΔX_0) in favor of one project has no effect on the standards and hence makes that project more likely to be chosen.

The remainder of the section is dedicated to proving the theorem. A common approach to solving a stochastic control problem like the one in (O) is to heuristically derive the HJB equation, establish the existence of a smooth solution to it, and appeal to a verification theorem that such a solution is in fact the principal's optimal value function. However, this direct approach has two limitations.

Given the finite horizon, the HJB equation would be a partial differential equation, and so establishing the existence of a smooth solution to it is not straightforward. Moreover, even if its existence were guaranteed, establishing qualitative properties of the optimal selection rule would be hard without an explicit characterization of a solution to the PDE. So, we adopt a different route. Because the argument is not typical of the optimal contracting literature, we describe our approach in section 5.1 below. We then present the formal results.

5.1. Approach to characterizing the optimal selection rule

Our first technical step is to consider relaxations of the principal’s problem that allow for *weak solutions*.¹⁸ Recall that in the control problem (O), the principal has to choose interim expected project choice q and project sensitivities C . Allowing weak solutions means we now allow the principal to additionally choose the underlying Brownian motions that drive projects’ random evolution (while still respecting the *law* governing this evolution as stated in the model section).¹⁹ In a typical discrete-time model, such a relaxation would be irrelevant, but in the present setting, it is a useful tool for the analyst. Given this broader definition of a control, we then proceed to show restricting attention to controls that have various economically intuitive features is without loss of optimality.

In Lemma 1, we show the principal sets (c^Δ, c^Σ) to be somewhat large until either she decides on her choice of project or the deadline arrives. Notice from (3) that q_t is unchanged during a phase when $c^\Delta = 0$. Thus, conjecturing that the principal cannot be worse off if she backloads collaboration to a time when it may be better targeted is reasonable. The lemma establishes this intuition applies more generally—the principal can backload collaboration by continuously speeding up decision-making whenever $\|(c^\Delta, c^\Sigma)\|$ is too small, thus creating residual time at the end for collaboration.

In Lemma 2, we show the principal ignores aggregate shocks. Given our earlier observation that c^Σ does not affect the principal’s objective function or agent

¹⁸Although 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, optimality of such a contract then follows.

¹⁹Other papers on continuous-time optimal contracting, and elsewhere in the optimal control literature, appeal to weak solutions of stochastic differential equations (e.g., Sannikov, 2008). We believe our specific use of weak solutions is novel. Namely, we solve a *relaxed program* entailing weak solutions as a solution method to characterize principal-optimal IC strong solutions. We see this difference as significant because strong solutions are arguably more natural for organizational design settings in which the shock process is beyond the principal’s control.

incentives, and that aggregate shocks are not an informative signal of agents' choices, it is intuitive that the principal should set $c^\Sigma = 0$.

In Lemma 3, 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 that moment, until the deadline. In particular, it is optimal for the principal to have the agents stop competing at some time, make a constrained-efficient choice with the partial information she has, and switch to collaboration on the chosen project thereafter.

The result from Lemma 3 is non-trivial. But once we have established competition is ended once and for all, showing the threshold for doing so should decline over time is straightforward. We formalize this fact in Lemma 4, by showing the stopping rule is a decreasing threshold²⁰: 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 $\mathcal{C} = \langle \Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P}, B, C, q \rangle$ 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 an \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 $|q_t| < 1$.

²⁰Studying costly sequential sampling problems that a single decision-maker faces, [Fudenberg et al. \(2018\)](#) show decreasing threshold rules can arise even without a deadline.

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.

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

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 $\tau_{\mathcal{C}}$.

Note 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. Decide quickly

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

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

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

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 “slowing down the clock” without changing the trajectory of the expected project choice, thus simply speeding up the original decision-making process and creating some residual time at the end. The benefit of rescaling time in this way is that this “extra” residual time can be utilized for efficient collaboration on a chosen project for a flow benefit of 1. Of course, the cost of this speeding up is that the duration for collecting flow payoffs is reduced. Note that holding 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. Thus, the cost of lost flow payoff as a result of speeding up is always less than the benefit of the extra collaboration time.

5.4. Respond only to relative performance

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

LEMMA 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 $|q_t| < 1$.*

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

For the interested reader, the proof of this lemma is a good example of why weak solutions are especially useful. The conclusion of Lemma 2—that c_t^Σ is almost surely zero—is a natural conjecture because c_t^Σ neither affects players’ incentives

nor enters the objective function. But establishing this conjecture formally (by standard methods) would require us to write down the HJB equation corresponding to (O) and prove both existence of a solution and concavity. The difficulty is that with a discontinuous flow payoff, our HJB equation does not belong to a class of well-understood partial differential equations. Reasoning through weak solutions circumvents this challenge. To see the basic idea, suppose that, on some paths, we have $c_t^\Sigma \geq 1$ and $c_t^\Delta = 0$. Then, we would conjecture that swapping c_t^Σ and c_t^Δ on those histories would lead to a payoff improvement (flow payoff being $|q_t| \leq 1 \leq c_t^\Sigma$) while keeping the law unchanged. But such a swapping operation changes the path of q_t . Permitting weak solutions enables us to construct new Brownian motions using the original ones and to work with these alternate controls that deliver a payoff improvement while preserving the law of motion. Although swapping arguments of this kind have been commonly used in studying dynamic incentive problems in discrete time, we hope our toolkit can facilitate similar arguments in continuous-time models.

5.5. First compete, and then collaborate on the winner

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

LEMMA 3: *For any 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. This lemma reduces our principal's problem to a non-strategic optimal stopping problem. To see why, notice $\{Y_t = \frac{1}{2}(\Delta X_0 + \Delta B_t)\}_t$ is an exogenous Brownian motion. Subtracting the constant term T from the principal's objective, the lemma says her modified objective is no greater than

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

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

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

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

LEMMA 4: *A function $\bar{z} : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ exists, such that for any Brownian base \mathcal{B} and any $(T, z) \in \mathbb{R}_+ \times \mathbb{R}$, $\tau_{T,z,\mathcal{B}}^* := \inf\{t \in [0, T] : |z + \Delta B_t| \geq \bar{z}_{T-t}\}$ is a (\mathcal{B}, T) -stopping time,²¹ 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 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. Thus 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.

²¹A (\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.

5.7. Characterizing the optimal selection rule

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

Proof of Theorem 1. Taking \bar{z} to be the function delivered by Lemma 4, let Π^* be the principal value generated by the behavior named in the theorem. Note 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 the lemmas referenced above. Lemma 2 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 3 then tells us the stopping time $\hat{\tau} := \tau_{\hat{\mathcal{C}}}$ 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 4 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 $\tau_{T,z,\hat{\mathcal{C}}}^*$.

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.²² 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.²³ Nevertheless, because collaboration starts early in expectation, the probability of an error approaches zero as the project horizon grows long.

6.2. Cancellation of projects before the deadline

In our setting, the principal chooses an optimal stopping time at which she makes a choice and then has both agents collaborate on the chosen project. An alternative interpretation is one in which the principal chooses when to irreversibly cancel one of the projects, after which both agents must work on the remaining project. A richer contracting environment in which the principal can choose to irreversibly cancel is more consistent with our motivating applications and is indeed equivalent to our current model. On the one hand, the principal cannot be worse off in the richer environment, because she can always abstain from canceling projects. Conversely, the principal can always simulate cancellation through a selection rule by deciding on a project in advance and having the agents collaborate on the chosen project. Moreover, in the richer model allowing irreversible termination, our optimal selection rule from Theorem 1 can be implemented in equilibrium, *without commitment*. The principal could simply terminate the project that is lagging behind by the current lead threshold, with each agent working on their own project unless it is canceled.

6.3. Agent indifference in the collaboration phase

Under our optimal rule, when agents collaborate, they are indifferent between competing and collaborating. Such indifference is common in many standard

²²That is, some finite-mean random variable τ_∞ exists such that the duration of the competition phase is first-order-stochastically dominated by τ_∞ . Indeed, one could take τ_∞ to be the optimal stopping time from an analogous stopping problem with no deadline, which is known to exhibit a constant lead threshold—the proof of Lemma 4 notes that it is $\frac{1}{\sqrt{2}}$. The constant threshold is finite because the option value of continuing vanishes with the probability of the Brownian motion revisiting zero, and a finite $|\Delta X|$ threshold is surpassed in finite expected time because the constant volatility is non-zero.

²³This observation is of course a consequence of the projects’ evolution being a Brownian motion and thus having unbounded supports.

contracting environments, but constructing a similar contract with strict incentives is typically possible by offering a small positive transfer to agents.²⁴ In our setting with no transfers, no obvious way exists to turn weak incentives into strict ones. Indeed, the best that the principal can achieve in any strict equilibrium is simply the no-commitment solution.²⁵

But, following the spirit of classical contract theory arguments, we ask what happens if the principal could provide monetary incentives up to a fixed budget of $\epsilon > 0$ (assuming for simplicity that agents have separable preferences over money and project choice). Observe that, in this case, implementing the optimum from Theorem 1 in a strict equilibrium would be easy. 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 collaborate on the chosen project, and it does not distort incentives in the initial phase of competition, 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 $\epsilon \rightarrow 0$. One can also similarly modify the model with small-scale monetary incentives to accommodate a small effort cost or a small preference for an agent to work on their own project.

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

²⁴For 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.

²⁵Given 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.

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. [1](#)
- 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. [7](#)
- J.-M. Benkert and I. Letina. Designing dynamic research contests. *American economic journal: Microeconomics*, 12(4):270–89, 2020. [6](#)
- J. Birkinshaw. Strategies for managing internal competition. *California Management Review*, 44(1), Fall 2001. [1](#)
- P. Bolton and C. Harris. Strategic experimentation. *Econometrica*, 67(2):349–374, 1999. [6](#)
- A. Bonatti and J. Hörner. Collaborating. *American Economic Review*, 101(2): 632–63, 2011. [6](#)
- A. Bonatti and H. Rantakari. The politics of compromise. *American Economic Review*, 106(2):229–59, 2016. [6](#)
- A. N. Borodin and P. Salminen. *Handbook of Brownian motion-facts and formulae*. Birkhäuser, 2012. [34](#)
- S. Callander and B. Harstad. Experimentation in federal systems. *The Quarterly Journal of Economics*, 130(2):951–1002, 2015. [6](#)
- J. M. Carey and J. Polga-Hecimovich. Primary elections and candidate strength in latin america. *The Journal of Politics*, 68(3):530–543, 2006. [1](#)
- R. M. Cyert and J. G. March. A behavioral theory of the firm. *Englewood Cliffs, NJ*, 2(4):169–187, 1963. [5](#)
- R. Deb, M. M. Pai, and M. Said. Evaluating strategic forecasters. *American Economic Review*, 108(10):3057–3103, 2018. [7](#)
- 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. [7](#)

- J. Farrell and T. Simcoe. Choosing the rules for consensus standardization. *The RAND Journal of Economics*, 43(2):235–252, 2012. 6
- D. Fudenberg, P. Strack, and T. Strzalecki. Speed, accuracy, and the optimal timing of choices. *American Economic Review*, 108(12):3651–84, 2018. 18
- R. Gibbons. March-ing toward organizational economics. *Industrial and Corporate Change*, 29(1):89–94, 2020. 6, 9
- Y. Guo and J. Hörner. Dynamic allocation without money. *Northwestern University and Yale University*, 2020. 7
- M. Halac, N. Kartik, and Q. Liu. Contests for experimentation. *Journal of Political Economy*, 125(5):1523–1569, 2017. 6
- A. V. Hirsch and K. W. Shotts. Competitive policy development. *American Economic Review*, 105(4):1646–64, 2015. 6
- B. Holmström. Moral hazard in teams. *The Bell Journal of Economics*, pages 324–340, 1982. 6
- 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. 11, 29
- I. Karatzas and S. E. Shreve. Brownian motion. In *Brownian Motion and Stochastic Calculus*, pages 47–127. Springer, 1998. 13, 29
- G. Keller, S. Rady, and M. Cripps. Strategic experimentation with exponential bandits. *Econometrica*, 73(1):39–68, 2005. 6
- P. Legros and H. Matsushima. Efficiency in partnerships. *Journal of Economic Theory*, 55(2):296–322, 1991. 6
- J. G. March. The business firm as a political coalition. *The Journal of politics*, 24(4):662–678, 1962. 5, 9
- A. M. Marino and J. Zabochnik. Internal competition for corporate resources and incentives in teams. *RAND Journal of Economics*, pages 710–727, 2004. 6
- A. McClellan. Experimentation and approval mechanisms. *Unpublished working paper*, 2021. 7
- M. A. Meyer. Learning from coarse information: Biased contests and career profiles. *The Review of Economic Studies*, 58(1):15–41, 1991. 7

- M. A. Meyer. Biased contests and moral hazard: Implications for career profiles. *Annales d'Economie et de Statistique*, pages 165–187, 1992. [7](#)
- D. Mookherjee. Optimal incentive schemes with many agents. *The Review of Economic Studies*, 51(3):433–446, 1984. [6](#)
- G. Moscarini and L. Smith. Optimal dynamic contests. *Science (June)*, pages 1–18, 2007. [6](#)
- G. Peskir and A. Shiryaev. *Optimal stopping and free-boundary problems*. Springer, 2006. [35](#), [36](#)
- T. J. Peters and R. H. Waterman Jr. *In Search of Excellence: Lessons from America's Best-Run Companies*. HarperCollins, 2003. [1](#), [2](#), [4](#)
- D. Ryvkin. To fight or to give up? dynamic contests with a deadline. *Management Science*, 68(11):8144–8165, 2022. [6](#)
- Y. Sannikov. A continuous-time version of the principal-agent problem. *The Review of Economic Studies*, 75(3):957–984, 2008. [5](#), [7](#), [17](#)
- N. L. Stokey. Wait-and-see: Investment options under policy uncertainty. *Review of Economic Dynamics*, 21:246–265, 2016. [4](#)
- T. Yamada. On a comparison theorem for solutions of stochastic differential equations and its applications. *J. Math. Kyoto Univ.*, 13(3):497–512, 1973. [29](#)

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

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

admits a unique strong solution. Optimality then follows readily from a comparison theorem. Indeed, following identically the proof of Theorem 2.1 in [Ikeda and Watanabe \(1977\)](#), any alternative control has a (weakly) first-order-stochastically dominated distribution of $|\Delta X_T|$.²⁶ \square

7.2. Proof of Lemma 1

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

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

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

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

²⁶That result shows a control $\Delta a_t = -\operatorname{sign}(\Delta X_t)$ minimizes $|\Delta X_T|$ —in fact, minimizes each of $\{|\Delta X_t|\}_{t \in [0, T]}$ —in an FOSD sense. However, reproducing the proof nearly verbatim establishes that control $\Delta a_t = \operatorname{sign}(\Delta X_t)$ maximizes $|\Delta X_T|$.

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

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

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

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

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

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

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

variation is zero. And indeed,

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

□

7.3. Proof of Lemma 2

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

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

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

7.4. Proof of Lemma 3

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

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

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

7.5. Proof of Lemma 4

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

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

Say a (\mathcal{B}, T) -stopping time is **optimal (given (\mathcal{B}, T))** if it maximizes $\mathbb{E} \left[\frac{1}{2} |z + \Delta B_\tau| - \tau \right]$ 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 1: 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 2: *A continuous function $v : \mathbb{R}_+ \times \mathbb{R} \rightarrow \mathbb{R}$ exists, such that for any Brownian base \mathcal{B} and any $(T, z) \in \mathbb{R}_+ \times \mathbb{R}$, the (\mathcal{B}, T) -stopping time*

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

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 1.

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,

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

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 1 and continuity of $v_{\mathcal{B}}$, Corollary 2.9 from [Peskir and Shiryaev \(2006\)](#) implies $\tau_{T,z,v_{\mathcal{B}},\mathcal{B}}$ is optimal, thereby generating $\mathbb{E} \left[\frac{1}{2}|z + \Delta B_{\tau_{T,z,v_{\mathcal{B}},\mathcal{B}}}| - \tau_{T,z,v_{\mathcal{B}},\mathcal{B}} \right] = v_{\mathcal{B}}(T, z)$. 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 4. Let $v : \mathbb{R}_+ \times \mathbb{R} \rightarrow \mathbb{R}$ be as delivered by Lemma 2, and define the set $G := \{(T, z) \in \mathbb{R}_+ \times \mathbb{R} : v(T, z) > \frac{1}{2}|z|\}$, which is relatively open in $\mathbb{R}_+ \times \mathbb{R}$ because v is continuous. For each $T \in \mathbb{R}_+$, let $G_T := \{z \in \mathbb{R} : (T, z) \in G\}$, which is open because G is. Let us make some easy starting observations about this family of sets. First, clearly, $G_0 = \emptyset$. Next, the set G_T is weakly increasing (with respect to set containment) in $T \in \mathbb{R}_+$. Indeed, v is nondecreasing in its first argument because, for any Brownian base \mathcal{B} and pair of times $t, T \in \mathbb{R}_+$ with $t \leq T$, every (\mathcal{B}, t) -stopping time is a (\mathcal{B}, T) -stopping time too. Finally, each G_T is symmetric about 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 $\mathbb{E} [\frac{1}{2}|z + \Delta B_\tau| - \tau]$ 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 $\mathbb{R} \setminus G_T \supseteq (-\infty, -\frac{1}{\sqrt{2}}] \cup [\frac{1}{\sqrt{2}}, \infty)$, the property would follow if we knew both $\mathbb{R}_+ \setminus G_T$ and $\mathbb{R}_- \setminus G_T$ were convex. 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

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

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 2, and let $\bar{v} := \max v([T, T+1] \times \{z - \epsilon, z + \epsilon\}) \in \mathbb{R}$. Then, any $s \in (0, 1]$ has

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

Observe now that $\tau < s$ if and only if the absolute value of 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 $|W_s| > \frac{\epsilon}{\sqrt{2}}$,²⁷ 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,

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

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 $\mathbb{E} \left[\frac{1}{2}|z + \Delta B_\tau| - \tau \right]$. Letting $\tau^* := \tau_{T, z, \mathcal{B}}^*$, Claim 2 establishes that $\tau \geq \tau^*$ almost surely. Assume now, for a contradiction, that τ is not almost surely equal to τ^* . Let us observe that some (\mathcal{B}, T) -stopping time $\tilde{\tau} \leq \tau$ exists such that, with positive probability, $\tau > \tilde{\tau}$ and $|z + \Delta B_{\tilde{\tau}}| > \bar{z}_{T-\tilde{\tau}}$.²⁸ But then, defining the

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

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

alternative (\mathcal{B}, T) -stopping times

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

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

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

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