The Dynamics of Multi-Project Collaborations

Charles Angelucci and Roi Orzach*

March 19, 2024

Abstract

We model multi-project collaborative dynamics with self-enforcing incentives. Two players collaborate across several domains, each offering infinitely many ex-ante identical projects. Every period, they jointly explore or exploit projects in each domain and make transfers. After exploring a project, players learn its benefits, which are asymmetrically distributed. We show that common features of collaborations such as gradualism, lengthy exploration, the postponement of project exploitation, the engagement in temporary project exploitation, or the return to previously abandoned projects occur in equilibrium.

Keywords: Collective Experimentation, Relationship Building, Gradualism.
JEL Classification Numbers: D21, D70, D83, L25

*Angelucci: MIT Sloan School of Management, 50 Memorial Drive, Cambridge MA 02142 (e-mail: cangeluc@mit.edu); Orzach: Department of Economics, MIT, 50 Memorial Drive, Cambridge MA 02142 (e-mail: orzach@mit.edu).
We thank Abhijit Banerjee, Daniel Barron, Alexander Bloedel, Alessandro Bonatti, Renee Bowen, David Byrne, Yifan Dai, Joyee Deb, Mert Demirer, Wouter Dessein, Glenn Ellison, Dana Foarta, Guillaume Fréchette, George Giorgiadis, Robert Gibbons, Marina Halac, Johannes Horner, Amit Khandelwal, Danielle Li, Jin Li, Rocco Macchiavello, Chiara Margaria, Simone Meraglia, Aroon Narayanan, Juan Ortner, Hazhir Rahmandad, Daniel Rappaport, Patrick Rey, Kareen Rozen, Klaus M. Schmidt, Steven Tadelis, Michael Ting, Jean Tirole, Marta Troya Martinez, Michael Whinston, and seminar audiences at the Oliver E. Williamson Seminar at Berkeley Haas, LMU Munich, the MIT Organizational Economics and Theory groups, the 2023 Montreal IO Conference, the Fall 2023 NBER Organizational Economics Workshop, the Paris School of Economics, the 2023 SIOE Conference, and the TSE Microeconomic Theory Workshop.

1 Introduction

Collaborations between firms in research and development (R&D) play a pivotal role in driving business success, from buyer-supplier relationships and joint ventures to strategic alliances. These collaborations are intricate and require careful development, often involving extensive experimentation and the management of multiple joint projects. For instance, in buyer-supplier relationships, companies frequently come together to create an array of products tailored to meet the demands of different markets. Similarly, R&D alliances typically involve collaborations across distinct areas of technological advancement. As one example, in the pharmaceutical sector, research-focused companies might form R&D alliances to combine their resources across several areas like novel drug discovery, the development of breakthrough therapies, and the implementation of pioneering biotechnologies.

In these broad collaborations, firms face the challenge of making joint decisions regarding the exploration and exploitation of opportunities across all their domains of cooperation. These decisions are complicated by the need to ensure that each collaborating firm perceives continued participation in the partnership to be in their best interest. Often, imbalances emerge in the distribution of benefits, making some firms more advantaged than others in the collaboration. As noted in a McKinsey report on buyer-supplier relationships, "Some collaborations promise equal benefits for both parties. [...] In other cases, however, the collaboration might create as much value overall but the benefit could fall more to one partner than to the other." (Benavides et al., 2012). As a result, the parties deriving most benefits must incentivize cooperation from those gaining less by promising to share the benefits. However, the effectiveness of these promises hinges on their credibility, which depends on the overall net present value of the collaboration. This value, in turn, is shaped by the net present value of each cooperative domain, fostering inter-dependencies among domains, both within the present moment and across different time periods, that affect the firms' exploration and exploitation choices.

In this paper, we develop a model of relationship building that examines the dynamics of multi-project collaborations when incentives are self-enforcing. We show that, despite the parties' desire to create common value swiftly, the process of establishing a successful collaboration is marked by a gradual expansion in the scope of the relationship and extended experimentation. The model also predicts other phenomena such as the postponement of project exploitation, the engagement in temporary project exploitation, and the return to previously abandoned projects. In Section 6, we discuss our main findings in light of the applied literature on collaborative dynamics and the literature regarding persistent productivity differences across firms.

The key features of our model are as follows. We consider a discrete-time framework where two players interact repeatedly over an infinite time horizon, with all actions being publicly observable. The players have the opportunity to engage in collaboration across a fixed number of domains that are identical ex ante. Specifically, the players can opt to cooperate on one project per period within each domain, where each domain corresponds to an infinite pool of projects that are identical and independent ex ante. Working individually on projects is not an option for the players; a successful outcome on any given project requires both parties' collaboration. Within each domain, any project that has been chosen before can be chosen again for collaboration, a scenario we call "project exploitation." The benefits of each project are time-invariant but initially uncertain, and they may vary asymmetrically across the players. Cooperation on a new project, or "project exploration," immediately reveals that project's benefits. Moreover, all projects entail a constant fixed cost for the players, during both the exploration and exploitation phases. As a result, players might be reluctant to cooperate in exploring projects if they expect that their individual benefit will not exceed the cost, and they may similarly be reluctant to cooperate in exploiting a project if their realized individual benefit falls below the cost. Finally, players can transfer money to each other, but these transfers are voluntary.

We focus on relational contracts (i.e., Subgame Perfect Equilibria) that maximize the players' joint surplus. In the main setting, we assume that each project exclusively benefits a single player. To highlight the challenges inherent in building common value within such asymmetric contexts, we contrast the dynamics that emerge under the optimal relational contract against those from a benchmark scenario with symmetric benefits where, in equilibrium, the players can act as a single player maximizing joint-surplus.

In the first part of the analysis, we assume not only that every project benefits the same player, but also that it is known which of the two players receives the benefits. This known asymmetry in the distribution of benefits can be illustrative of a buyer-supplier collaboration, in which the buyer is the sole recipient of the revenues generated from the sale of the final product. In this initial analysis, to focus on the mechanisms through which the player who benefits from the projects succeeds in incentivizing the other player to collaborate in both the exploration and exploitation of projects, we further assume that all projects yield benefits that are either zero or a fixed constant. When the discount factor is high, the decisions of the players regarding the exploration and exploitation of projects mirror those observed in the symmetric-benefits benchmark. They immediately explore projects in all domains and, in each domain, exploit the first project they find that generates positive benefits, using monetary transfers to redistribute the benefits. When instead the discount factor is low, we show that the players adopt a gradual approach, expanding the scope of their relationship (defined by the number of domains within which they are either exploring or exploiting projects) incrementally over time (Propositions 3 and 4). By expanding scope gradually, the players are able to initiate their collaboration, with both the low initial costs for the non-benefiting player and the prospect of future scope expansions enabling the benefiting player to credibly offer compensation. Concurrently, the discovery of projects worth exploiting raises the value of the relationship and makes it possible for the players to expand the scope of their collaboration by exploring projects in new domains. Furthermore, when the initial value of the relationship is particularly low, the players might have no choice but to gradually extend the scope of their relationship by delaying the exploitation of valuable projects. This strategy lowers the immediate costs for the non-benefiting player, enabling the exploration of projects in new domains. Notably, when this approach is chosen, whether the players' relationship ever reaches its maximum potential scope can be uncertain (Proposition 5).

In the second part of the analysis, we assume that both players are equally likely to benefit from any given project. As a result, both players are motivated to collaborate in the exploration of projects, leading to a maximal relationship scope in all periods. However, for the exploitation phase of a project, one player must still provide incentives to the other. Moreover, we posit that the benefits derived from projects follow a distribution with convex support, so our analysis focuses on which projects the players choose to exploit. When the discount factor is high, the players' actions again replicate those that arise in the symmetric-benefits benchmark: they treat the exploration of projects independently across domains and opt for a project's exploitation when its benefits exceed a common threshold set for all projects. When the discount factor is low, the players' search for exploitable projects becomes interdependent. Specifically, the thresholds used for assessing whether a project is worth exploiting permanently are co-determined and differ across projects. This co-determination arises because selecting a project for exploitation impacts the overall value of the relationship and, consequently, the players' ability to cooperate on exploiting other projects.

We first show that the permanent exploitation of projects across all domains occurs only when each project is valuable enough that the players would exploit it under symmetric benefits, and further, only when the average value of the projects exceeds a threshold (Proposition 6). This additional condition guarantees that the relationship's value is sufficiently high, enabling the players to collaborate in the exploitation of each project. Moreover, it implies that the resulting time at which a project is permanently exploited is likely to substantially exceed that observed under the symmetric-benefits benchmark. Moreover, in their quest to identify projects for permanent exploitation, we show that the players may opt for the temporary exploitation of certain projects with the understanding that they might abandon them later, or they may bypass some projects, only to return to them when the value of their relationship has grown sufficiently to enable exploitation (Proposition 7). In sum, inefficiencies arise not only in terms of the time spent exploring projects, but also in the departure from the decision rule optimal for a single decision-maker consisting of either permanently exploiting projects or permanently abandoning them.

Our theoretical analysis ends with an extension of the model that includes projects that yield both symmetric and asymmetric benefits. We show that the players use symmetric projects as stepping stones, enabling cooperation in identifying more lucrative but asymmetric projects. We also find that the players set lower exploitation thresholds for symmetric projects compared to asymmetric ones. This difference occurs because symmetric projects do not require promises of compensation to enable exploitation.

Finally, we examine the applied literature on collaborative dynamics, interpreting it through the framework of our model. Patterns of gradualism and prolonged experimentation are commonly observed in buyer-supplier relationships. Moreover, the slow growth in relationship value helps explain the surprisingly robust nature of buyer-supplier relationships, a phenomenon documented within the trade literature. Additionally, we discuss how our model can be interpreted to capture internal firm dynamics, thereby contributing to the literature highlighting the role of managerial practices in driving persistent productivity differences.

The rest of the paper is organized as follows. Section 1.1 discusses the related theoretical literature. Section 2 describes the model. Section 3 characterizes the set of optimal relational contracts we focus on and analyzes a benchmark scenario with symmetric benefits. Section 4 solves the model, first by focusing on the dynamics of the scope of the players' relationship and then by analyzing the dynamics of the players' project exploitation choices. Section 5 analyzes extensions and Section 6 reviews the applied literature on collaborative dynamics. Section 7 concludes.

1.1 Related Theoretical Literature

In this section, we review the theoretical literature related to our work. We postpone the discussion of the applied literature to Section 6.

Firstly, our research connects to the large literature on multi-armed bandit problems, dating back to Weitzman (1979). For a review of applications within economics, see Bergemann and Välimäki (2008). A subset of this literature analyzes strategic interactions. Bolton and Harris (1999) and Keller et al. (2005) consider settings in which players independently pull arms and free-ride on each others' experimentation.[1] In Strulovici (2010), players collectively choose between a safe arm and a risky one, with its asymmetric benefits revealed over time through experimentation. Bonatti and Hörner (2011) examine a scenario in which a group of agents collaborates in a collective experimentation process, characterized by private effort choices. Further, Reshidi et al. (2021) and Chan et al. (2018) contrast group and individual decision-making regarding experimentation, looking at the impact of static versus sequential information acquisition and of voting rules. Similar to these papers, our focus is on strategic interactions. The main distinctions are that our setting is characterized by infinitely repeated interactions and heterogeneous preferences, includes transfers, and permits players to collectively experiment with multiple projects simultaneously. We analyze a general setting characterized by analytically tractable, yet rich, dynamics, and a wide array of applications.

Secondly, this work is related to the literature on relational contracts (see e.g., Bull, 1987; Macleod and Malcomson, 1989; Baker et al., 1994, 2002; Levin, 2003, for early contributions).[2] The positive feedback effect between the value of players' relationship

[1] For more recent work along these lines, see Anesi and Bowen (2021) and Hörner et al. (2022).

[2] Also at the intersection of the bandit and the relational contracting literatures, Urgun (2021) examines a scenario where a principal interacts with multiple agents whose publicly-observable types

and incentive strength is pervasive across relational contracting models. However, it rarely produces dynamics, because current production is typically influenced only by present actions, not by past choices as in our setting. An exception is Halac (2014), who studies a setting in which the value of the players' relationship increases with the duration of the relationship. The players initially choose to cooperate on low-risk, low-return projects, and they switch to high-risk, high-return projects once their relationship has grown sufficiently valuable.[3] In our setting, it is the discovery of projects worthy of exploitation that increases the value of the relationship. We analyze the implications of this effect on the players' choices between project exploration and exploitation when they are engaged in multiple projects. In contrast, Chassang (2010) analyzes a setting where increases in relationship value diminish the players' motivation to enhance their collaboration. In his model, the agent knows which arms are productive and which are not, while the principal, at the outset, cannot differentiate between the two. Without monetary incentives, incentivizing the agent to choose productive arms is accomplished by the threat of firing the agent following failures. This dynamic makes motivating exploration progressively expensive as more productive arms are identified. Should the relationship endure, it ultimately enters an "exploitation" phase and its value stops growing. In our model, the players are symmetrically informed about their environment, and the presence of transferable utility removes the need for inefficient on-path punishments. These two features lead to the positive feedback effect mentioned above.[4]

Our work also connects to Bernheim and Whinston (1990), who analyze firms operating in multiple markets, showing that maintaining collusion in easier markets can help support collusion in more challenging ones. Similarly, Levin (2002) shows the advantages firms gain by pooling heterogeneous employees' incentives into a "multilateral" relational contract. In our setting, the ability to pool relational incentives across multiple projects is key in generating gradualism in relationship scope.

Finally, we add to the body of research that examines gradualism. Watson (1999, 2002) examine a setting in which players are uncertain regarding their counterpart' in-

depend on the contracting history.

[3]In Halac (2015), a principal leverages this feedback effect by making an upfront and relationship-specific investment prior to her repeated interaction with an agent.

[4]Introducing transferable utility within Chassang (2010), where information asymmetry plays a central role, would make the value of the players' relationship constant on path. For a setting similar to Chassang (2010) but with imperfect transfers and uncertainty about the value of the relationship, see Venables (2013). For a bandit problem embedded in a principal-agent setting, see Ide (2024).

tentions—to either collaborate genuinely or take advantage of the other. The players begin with low cooperation to mitigate the losses from defection. As the players become more optimistic, the collaboration grows. Collaborations involving trustworthy players achieve optimal cooperation, while those with untrustworthy players eventually end. In our setting, the relationship develops incrementally, not due to screening intentions, but because credibility is built by the players over time.[5]

2 The Setup

We suppose that there are 2 players who have the opportunity to interact at different time periods $t = 0, 1, 2, \ldots$. Each player, denoted by $i = 1, 2$, has a discount factor δ and a per-period outside option equal to zero. The players' interaction spans m "domains," where m is fixed exogenously. For each domain $j = 1, \ldots, m$, there exists an infinite set of projects $\mathcal{P}_j$, where $\mathcal{P} = \cup_j \mathcal{P}_j$. In each period t, and for each domain j, the players may select up to one project from the set $\mathcal{P}_j$. We assume that the players cooperate on a project if and only if both players choose it, thus following a unanimity rule. We denote by P_i^t the finite set of projects chosen by player i in period t and by $\mathbf{P}^t$ the corresponding set of projects the players cooperate on, where $\mathbf{P}^t = P_1^t \cap P_2^t$. We refer to $|\mathbf{P}^t| \leq m$ as the players' "relationship scope" in period t.

Each project in $\mathbf{P}^t$ imposes a cost of $c > 0$ on both players. Further, each project $p \in \mathcal{P}$ is associated with a vector of initially unknown time-invariant individual valuations $(v_{p,1}, v_{p,2}) \in \mathbf{R}_+^2$. A project's associated individual valuations are publicly observed immediately after the players have cooperated on it for the first time. We say that a project p is being "explored" in period t if the players cooperate on it for the first time, and that it is being "exploited" during period t if the players have chosen to cooperate on it already in some prior period. Note that we place no intertemporal restrictions on the set of projects that are available; for instance, nothing prevents the players from exploring a project, potentially exploiting it for several periods, then temporarily abandoning it, and returning to it at a later time. We refer to a project's

[5]Gradualism also arises in Ghosh and Ray (1996) and Kranton (1996), where players are randomly matched and can exit relationships at any time, with new partners possessing only limited information about the player's past history. In the context of corporate finance, investment levels can increase over time as borrowers gradually build collateral (see Tirole, 2006, Chapter 4 and references therein). In our scenario, the continuation value of the relationship functions like a collateral. However, this collateral is pledged by both players involved, rather than a single player.

sum of individual valuations as the project's value, and denote it by s_p. We assume that each project p's vector of individual valuations is distributed independently and identically across projects and domains. This assumption implies that all domains of the relationship are ex ante identical.

The players can exchange money twice during each period. At the beginning of each period t, the players make discretionary transfers to each other, where $w^t_{i,-i} \in \mathbf{R}^+$ denotes such a transfer from player i to player $-i$. At the end of each period t, players again make discretionary transfers to each other, where $b^t_{i,-i} \in \mathbf{R}^+$ denotes such a transfer from player i to player $-i$.[6]

Further, player i's period t payoff is equal to:

$$\pi^t_i = w^t_{-i,i} - w^t_{i,-i} + b^t_{-i,i} - b^t_{i,-i} + \sum_{p \in \mathbf{P}^t} (v_{p,i} - c), \text{ where } i, \in \{1,2\}. \tag{1}$$

Equation (1) implies that a key friction between the players will be about the choice of projects that benefit one player but not the other, and that money will serve the purpose of aligning incentives. As we will show, this friction will lead to complementarities across domains despite the absence of technological inter-dependencies.

We conclude the model's description by stating the timing of the stage game. Both players simultaneously choose their discretionary transfers $w^t_{i,-i}$. Next, both players simultaneously make their project choices P^t_i. For each project $p \in \mathbf{P}^t$, the players observe the vector $(v_{p,1}, v_{p,2})$, pocket their individual valuation, and incur the cost c. Finally, both players simultaneously choose their discretionary transfers $b^t_{i,-i}$.

Relational Contracts. A relational contract is a complete plan for the relationship. Let $h^t = (\mathbf{w}^0, \mathbf{P}^0, \mathbf{v}^0, \mathbf{b}^0, \ldots, \ldots, \mathbf{w}^{t-1}, \mathbf{P}^{t-1}, \mathbf{v}^{t-1}, \mathbf{b}^{t-1})$ denote the history up to date t and $\mathcal{H}^t$ the set of possible date t histories, where boldface lowercase letters indicate vectors. Then, for each date t and every history $h^t \in \mathcal{H}^t$, a relational contract describes: (i) the $\mathbf{w}^t$ transfers; (ii) the set of projects $\mathbf{P}^t(\mathbf{w}^t)$ as a function of $\mathbf{w}^t$; and (iii) the $\mathbf{b}^t(\mathbf{w}^t, \mathbf{P}^t, \mathbf{v}^t)$ transfers as a function of $\mathbf{w}^t$, $\mathbf{P}^t$, and the realizations of $\mathbf{v}^t$. Such a relational contract is self-enforcing if it describes a Subgame Perfect

[6] We incorporate the option of monetary transfers both before and after the players' project choices, although removing either would not qualitatively affect our results. Without transfers at the beginning of each period, surplus might no longer be fully redistributed across the players without affecting incentives. Without transfers at the end of each period, incentives for the current period would rely on transfers from the subsequent period, complicating the proofs.

Equilibrium of the repeated game. Within the class of Subgame Perfect Equilibria, we analyze pure-strategy equilibria which maximize the players' joint surplus.[7] In the event of a deviation in some period t, the players respond (i) by choosing $P_i^t = \emptyset$ and $b_{i,-i}^t = 0$ if these choices have not been made yet and (ii) by permanently breaking off their relationship (i.e., reverting to the worst equilibrium of the stage game from the next period onward). This punishment is without loss of generality as it occurs only out-of-equilibrium (c.f. Abreu, 1986).[8]

Examples. We present two concrete illustrations of the setting. In the first example, a buyer and a supplier engage in collaboration to manufacture a final product intended for distribution across various geographical markets. These markets each present unique local conditions, demanding tailored product customization. The experimentation process requires both firms to undertake non-contractible investments, such as worker training and marketing efforts. Even after experimentation concludes, ongoing investments will still be needed. Since all revenue generated from the sale of the final products accrues to the buyer, they must offer credible compensation to the supplier to guarantee their continued collaboration. In the second example, two firms establish an R&D alliance with the aim of pooling resources across various domains of research. To identify profitable collaborations within each domain, cooperation necessitates extensive experimentation efforts, involving repeated investments that cannot be fully formalized in contracts. Furthermore, the collaborations the firms settle on may result in uneven benefits across the firms, possibly due to the utilization of distinct strengths or resources in ways that were uncertain at the beginning. The firm reaping the greatest benefits from the collaboration might have to promise compensation to the other to secure their continued involvement.

[7] Restricting attention to pure strategy equilibria is without loss because (i) mixing on transfers cannot benefit the players as it increases the maximal transfers they promise each other and (ii) mixing on projects either results in inefficiently low relationship scope due to miscoordination or a limited relationship scope that can be replicated by players not choosing projects in some domains.

[8] Equivalently, in the period following a deviation, players could transition to a continuation equilibrium in which everything remains the same, except that the entire surplus is allocated to the player who did not deviate. This punishment provides identical incentives. Because it is Pareto optimal, it is also less prone to renegotiation.

3 Preliminary Analysis

In this section, we characterize the set of surplus-maximizing relational contracts our analysis focuses on. We also analyze the benchmark case where every project equally benefits both players.

3.1 Characterization of Optimal Relational Contracts

In our setting, surplus-maximizing relational contracts will depend on the players' beliefs about the projects. Denote by $\mu^t(h^t) := \{\Delta(v_{p,1}, v_{p,2})|h^t)\}_{p\in\mathcal{P}}$ the beliefs the players hold about the projects' valuations given all the observed valuations up through period $t-1$. We now show that there exist surplus-maximizing relational contracts that condition on h^t only through the beliefs $\mu^t(h^t)$. Formally, restricting attention to relational contracts that specify the same continuation equilibrium following any two on-path histories h_1^t and $h_2^{t'}$ that lead to the same beliefs μ is without loss. Further, the continuation equilibrium the relational contract prescribes is surplus-maximizing, in the sense that there does not exist another continuation equilibrium that generates a higher total surplus across the players. The following proposition formalizes this result and provides a necessary and sufficient condition for a given project selection rule (i.e., a mapping from beliefs to projects) to be implemented by a relational contract.

Proposition 1 (**Optimal Relational Contracts**)

- *For any surplus-maximizing relational contract, there exists an alternative surplus-equivalent relational contract such that (i) for all t and for all on-path histories $h^t \in \mathcal{H}^t$, the continuation equilibrium is surplus maximizing, and (ii) for any two on-path histories h_1^t and $h_2^{t'}$, if $\mu^t\left(h_1^t\right) = \mu^{t'}\left(h_2^{t'}\right)$, then the relational contract specifies the same continuation equilibrium following these histories.*

- *There exists a relational contract that implements a project selection rule $\mathbf{P}(\cdot)$ if and only if the following inequality holds for all t and for all histories $h^t \in \mathcal{H}^t$:*

$$\sum_{p\in\mathbf{P}(\mu^t)} \sum_{i=1}^{2} \max\left(0, c - \mathbb{E}\left(v_{p,i}|\mu^t\right)\right) \leq \mathcal{C}\left(\mu^t\right), \tag{2}$$

where $\mathcal{C}(\mu^t)$ ("the continuation value") is the expected net present value of the players' relationship starting in $t+1$ given $\mathbf{P}(\cdot)$ and μ^t.

The proof of this proposition extends the work of Levin (2003) and is provided in the Online Appendix. In our setting, the players' continuation value is stochastic. We show that considering the expectation of the continuation value is sufficient to characterize which project selection rules can be implemented by a relational contract.

The intuition for the first statement of the proposition is based on the following two observations. First, any surplus-maximizing relational contract is necessarily surplus-maximizing following any on-path history, for otherwise non-surplus-maximizing continuation equilibria could be replaced with surplus-maximizing ones, with transfers appropriately designed to maintain all players' incentives. Second, when confining our attention to surplus-maximizing continuation equilibria, the only history-dependent outcome that can alter the set of optimal continuation equilibria are the players' beliefs μ^t about the projects. We call such relational contracts optimal.

Next, recall that the main tension faced by the players is that the project selection rule which maximizes their joint surplus may involve the selection of projects that do not benefit each player. Inequality (2) states that for a relational contract to implement a project selection rule everywhere on path, the continuation value must exceed the total reneging temptation across players and projects in all periods and for all possible histories. The total reneging temptation is the sum across players and across projects of a project's reneging temptation to a player. The sum is across projects because each player can deviate from the relational contract by selecting any subset of $\mathbf{P}^t$. In turn, a project's reneging temptation to a player is either equal to zero, in case the project generates a positive net expected gain to the player, or equal to the magnitude of the net expected loss. That the relational contract creates more continuation value to the players than the sum of their gains from defecting is necessary for the relational contract to constitute an equilibrium. In the proof, we show that the presence of money also ensures that this condition is sufficient.

Finally, we note that the second statement of the proposition means that characterizing the optimal relational contract can be reduced to characterizing the players' optimal project selection rule, which will thus be the focus of our analysis hereafter. To understand this, observe that all transfers cancel each other in the expression for the joint surplus of the players, as well as on the right-hand side of Equation (2).

3.2 Benchmark with Symmetric Benefits

We now analyze the benchmark case where every project benefits the players equally, ensuring that they have perfectly aligned incentives. Specifically, we suppose that for each project $p \in \mathcal{P}$, $v_{p,1} = v_{p,2}$.[9] The predictions this benchmark analysis produces are identical to those that would result if a single decision-maker, whose payoff is given by the sum of the payoffs of both players, were to make all the decisions.

Proposition 2 (**Symmetric Benefits Benchmark**)
When projects generate symmetric benefits, all optimal relational contracts specify a project selection rule that is identical and independent across all m domains of the players' relationship. Further,

1. *If the players select any project in any period t in domain j, then the players select a project in domain j in all periods.*
2. *There exists an increasing function $s^0(\delta)$ such that the players exploit project p in domain j if and only if $s_p \geq s^0(\delta)$.*

When projects yield equal benefits for both players, Inequality (2) from Proposition 1 simply states that the net present value of the net payoff resulting from the selection of any project (accounting for the potential abandonment of a project) must be non-negative. Because this feature will always hold under any optimal relational contract, Inequality (2) can be ignored. The intuition behind the players treating each domain of their relationship separately and identically follows from our assumptions wherein (i) payoffs are additively separable across projects, meaning there are no inter-dependencies like economies or diseconomies of scope, and (ii) all projects benefit the players equally. The intuition for statement (1) is that if the players find it rational to explore project $p \in \mathcal{P}_j$ in some period t, then exploration must exhibit a positive net present value of the players' net payoffs, accounting for the possibility of project abandonment. Since the players have access to infinitely many ex ante identical projects, they would thus always opt for exploration rather than the non-selection of a project and, by extension, a project will be selected in every period.

To gain intuition for statement (2), note that when the players find it optimal to select a project, the players either (i) exploit a previously-explored project p or (ii)

[9] The findings of this subsection remain valid for asymmetric benefits as long as each project either positively or negatively affects both players.

explore new projects with the hope of ultimately settling on a superior project in the future. The players never return to an abandoned project because we have assumed an infinite supply of ex ante identical projects. Finally, statement (2) also states that as the discount factor increases, the value of exploring alternative projects rises, since any superior project identified can be used across all future periods.

In sum, when benefits are symmetric, the players maximize the scope of their relationship at all times and they switch to permanently exploiting projects based on an independent, identical, and time-invariant threshold. These features will not always be true when benefits are asymmetric, a scenario we now analyze.

4 Analysis

We divide the analysis of the asymmetric benefits case into two subsections. In subsection 4.1, we assume that one player requires incentives to explore projects. Within this context, we analyze the dynamics of the scope of the players' relationship. In subsection 4.2, we assume that both players are motivated to explore projects and instead focus on their decision-making process regarding project exploitation. Throughout the analysis, we make the following assumption:

Assumption 1 (**Feasibility**) *With positive probability,* $s_p > \tilde{s}(\delta) := c\frac{1+\delta}{\delta}$.

By Proposition 1, Assumption 1 implies that there always exist projects that the players can cooperate in exploiting, despite the asymmetry of benefits across players. We denote the minimum value of s that fulfills Assumption 1 by $\tilde{s}(\delta)$.

4.1 The Dynamics of Relationship Scope

We specialize the model as follows. We assume that, $\forall p$, $s_p \in \{0, v\}$, where $v > 2c/q$. Specifically, $s_p = v$ with probability q, independently across projects. These assumptions streamline our focus on the players' relationship scope by simplifying their decision-making: for instance, players will not exploit projects valued at 0 and will not explore in domains where a project with value v is already found.[10] We also assume player 2 gains no benefit from any project, meaning that $v_{2,p} = 0$ $\forall p$,

[10] The analysis remains qualitatively identical when considering distributions with convex support. Also, $qv > 2c$ implies that exploring domains is socially efficient independently of the discount factor.

and, therefore, must be incentivized to both explore and exploit projects by player 1. Throughout, we refer to projects with value v as "suitable projects."

By Proposition 1, the optimal relational contract depends only on the number of suitable projects identified to date. Denote this number by n and denote (i) by $f_{\text{explore}}(n)$ the number of domains where the players are exploring projects, (ii) by $f_{\text{exploit}}(n)$ the number of projects the players are exploiting, and (iii) by $f(n) = f_{\text{explore}}(n) + f_{\text{exploit}}(n)$ the corresponding scope of the players' relationship. The following proposition provides key properties of the function $f(n)$.

Proposition 3 (**Scope**)
To any (non-empty) optimal relational contract corresponds a function $f(n)$ specifying the scope of the players' relationship. Such a function satisfies the following two conditions:

1. *$f(n)$ is monotonically increasing in n for all n.*
2. *$f(n) \geq \min\left(\lfloor \frac{n}{c}\frac{\delta}{1-\delta}(v-2c)\rfloor, m\right) \geq n$ for all n.*

Proof of Proposition 3. Statement 1: Any project selection rule that the players implement when they have identified $n-1$ suitable projects remains implementable when they have identified n suitable projects. This occurs because the continuation value is weakly increasing in the number of suitable projects identified by the players.

Statement 2: By Assumption 1, the players can always exploit the n suitable projects they have identified to date. Further, the number of suitable projects discovered by the players can only increase over time. Thus, their continuation value when they have identified n suitable projects exceeds $n\delta(v-2c)\ /\ (1-\delta)$. Next, because exploring projects is always valuable when feasible (since $qv > 2c$), under any optimal relational contract the scope of the players' relationship is as large as feasible. In turn, a given scope is feasible if (i) it is an integer, (ii) if it is lower than m, and (iii) if the integer multiplied by c is lower than or equal to the continuation value (see Inequality (2) in Proposition 1). Combining these three conditions generates the first inequality. The second inequality follows from Assumption 1. □

Proposition 3 states that any optimal relational contract exhibits weak growth in the scope of the players' relationship along the equilibrium path. However, it does not rule out the possibility that the players always start with maximal scope, similar to the symmetric-benefits scenario considered in Proposition 2. We now show that, with asymmetric benefits, there exists a value of δ low enough such that both $0 < f(0) < m$ and the scope of the relationship increases over time.

Proposition 4 (Gradualism)
For any $m \geq 2$, there exist two thresholds $0 < \delta^* < \bar{\delta} < 1$. If $\delta \geq \bar{\delta}$, any optimal relational contract is such that the relationship scope is always maximal on path (i.e., $|\mathbf{P}^t| = m$ for all t). If $\delta \in [\delta^*, \bar{\delta})$, any optimal relational contract is such that the relationship scope is initially limited (i.e., $0 < |\mathbf{P}^0| < m$) on path, with the scope (i.e., $|\mathbf{P}^t|$) increasing at least once along the equilibrium path.

In the intuition that follows, we consider only the histories in which the constraints implied by Proposition 1 are most binding. Appendix A provides proofs showing the necessity and sufficiency of these constraints.

First note that, if they can, the players adopt the project selection rule of the symmetric-benefits benchmark, immediately exploring projects in all domains of cooperation, since this approach maximizes their joint surplus. In this case, the most binding incentive constraint implied by Proposition 1 is:

$$m \cdot c \leq m \cdot \mathcal{C}(\text{exploration}), \tag{3}$$

where $\mathcal{C}(\text{exploration})$ denotes the per-domain continuation value associated with the current exploration of a project. As δ diminishes, $\mathcal{C}(\text{exploration})$ decreases to a point where Inequality (3) fails to hold. In other words, immediate cooperation in all domains is implementable only if the discount factor is relatively high ($\delta \geq \bar{\delta}$).

The drawback of starting exploration in all domains simultaneously is that it prevents players from pooling incentives across domains, as evidenced by (3) not depending on m. For this reason, when $\delta < \bar{\delta}$, project selection rules in which relationship scope is limited but constant over time are not implementable either. Instead, gradual project selection rules, in which players initially collaborate in a limited number of domains and later expand their cooperation, emerge as the possible optimal alternative. A gradual approach reduces the total cost borne by player 2 in the initial stages by postponing the start of advantageous cooperation in additional domains. Furthermore, when the scope of their collaboration increases, the players are able to leverage the value derived from the suitable projects they have identified.

To illustrate, suppose the players start by exploring projects in a single domain and, upon finding a suitable project, begin exploring projects in the $m - 1$ remaining domains while exploiting the first project. There exists a relational contract that implements this specific project selection rule if and only if:

$$c \leq \mathcal{C}(\text{exploration}) + (m-1) \cdot \mathcal{C}(\text{delayed exploration}), \tag{4}$$

$$m \cdot c \leq (m-1)\mathcal{C}(\text{exploration}) + \mathcal{C}(\text{exploitation}), \tag{5}$$

where $\mathcal{C}$(exploitation) denotes the per-domain continuation value associated with the permanent exploitation of a project, and $\mathcal{C}$(delayed exploration) denotes the per-domain continuation value associated with the exploration of a project that will begin following the discovery of a suitable project in the first domain of cooperation. In words, Inequality (4) ensures that the players are able to explore a project in one domain anticipating the future exploration of projects in the $m-1$ remaining domains. Similarly, Inequality (5) ensures that the players are willing to explore projects in $m-1$ domains while simultaneously exploiting one project.

The right-hand sides of Inequalities (4)-(5) are strictly increasing in δ. Further, (4) is strictly easier to satisfy than (3) since $\mathcal{C}$(delayed exploration) > 0, and (5) is strictly easier to satisfy than (3) since $\mathcal{C}$(exploitation) $> \mathcal{C}$(exploration). As a result, the gradual project selection rule described above is implementable across a broader range of discount factors when compared to the symmetric-benefits project selection rule, and, therefore, there are always gradual project selection rules available to the players when the discount factor reaches a level that permits their implementation, albeit not sufficient for achieving maximal scope right from the outset.[11,12]

Proposition 4 states that the scope of the players' collaboration will be limited at first and expand over time when the discount factor is intermediate ($\delta \in [\delta^*, \bar{\delta})$). However, it does not state that scope necessarily reaches its maximum potential of m.[13] Moreover, Proposition 3 establishes a lower bound on the scope of the rela-

[11] An implication of the previous proposition is that as m increases, the range of discount factors for which the optimal relational contract is nonempty also widens. Intuitively, a larger m offers the players a wider range of project selection rules and more chances to pool incentives across cooperation domains. Further discussion of this feature and related findings can be found in Section 5.1.

[12] A more gradual project selection rule may not always be feasible across a wider range of discount factors compared to a less gradual approach, as it delays valuable collaboration.

[13] Addressing this question by obtaining closed-form analytical solutions for $f_{\text{exploit}}(n)$ and $f_{\text{explore}}(n)$ is intractable for a general value of m. To illustrate this complexity, consider the case where $n = m-2$. Although Proposition 3 states that $f(m-2) \geq m-2$, there remain numerous potential values for f_{explore} and f_{exploit}. The players might choose to exploit the $m-2$ projects and explore zero, one, or two additional projects, or the players might exploit $m-3$ projects and exploit zero, one, or two additional projects. Moreover, the optimal project selection rule is discontinuous

tionship as a function of n, which in turn sets a lower bound for the growth of the relationship's scope. If the lower bound on $f(n)$ exceeds $n+1$ for all values of n, then relationship scope converges to m. This could happen if a low value of q leads to a sufficiently low continuation value that the optimal project selection rule is gradual, while a large value of v ensures that identifying suitable projects substantially increases the continuation value, causing large expansions in relationship scope. We now show that if this sufficient condition is not satisfied, optimal gradual project selection rules might not achieve maximal scope.

Definition 1 (**Maximal Equilibrium Scope**)
The broadest scope of the players' relationship achieved on the equilibrium path is said to be stochastically maximal (respectively, deterministically maximal) if it equals m with a probability of $0 < p < 1$ (or, in the deterministic case, $p = 1$).

We are now in a position to present and prove the following results.

Proposition 5 (**Bounded Relationships**)
For any non-empty optimal relational contract, the broadest scope of the players' relationship achieved on the equilibrium path is (i) either stochastically or deterministically maximal, (ii) deterministically maximal when $m = 2$, but (iii) is stochastically maximal for an open set of parameter values when $m \geq 3$.

The reason why the scope of the relationship must reach its maximum with positive probability is as follows. If scope is bounded by the value $n < m$, then, with positive probability, the players will first identify $n-1$ suitable projects and, subsequently, explore projects in exactly one domain. However, if players can explore projects after identifying $n-1$ suitable projects, it becomes even easier for them to also do so upon identifying n suitable projects. Therefore, the players would never limit the scope of their relationship to a value n lower than m with certainty. One can use such a result to show that the scope must reach m with positive probability.

Next, the proposition states that when $m = 2$, the broadest scope the players achieve is deterministically maximal. If it were stochastically maximal, $f(1)$ would equal 1. However, by Proposition 3, the condition $0 < f(0) \leq f(1)$ would also have to

in the model parameters. Finally, the choice of project selection rule at $n = m - 2$ determines the continuation value at $n = m - 3$, which, in turn, determines what project selection rule is feasible at $n = m - 3$ (and in all earlier periods). Solving for the corresponding Bellman Equation in closed form is intractable.

hold. Yet, $f(0) = f(1) = 1$ implies that scope never reaches its maximum potential of 2, which we have just shown cannot occur. In contrast, we find that the broadest scope of the players' relationship attained on path is not always guaranteed to achieve m when $m \geq 3$. In Appendix A we show this result by focusing on the case where $m = 3$. Applying the same logic as in the preceding paragraph, scope must exceed 1 for any non-empty relational contract. However, scope can be bounded by 2. For this to happen, the following three conditions must hold: (i) When $n = 1$, players explore projects in two domains, since the relationship must reach m with positive probability; (ii) To explore projects in two domains, players must postpone exploiting the first suitable project that they found. This is necessary to prevent a situation where $f(1) > f(2)$, which has also been shown to be impossible; and (iii) When $n = 0$, the players explore projects in a single domain. Given the structure of the problem, one can show that if $f(0) = 2$, then $f(2) > 2$. The only project selection rule that satisfies these conditions is:

$$f_{\text{explore}}(n) = \begin{cases} n+1 & \text{if } n \leq 1 \\ 0 & \text{if } n > 1 \end{cases} \qquad f_{\text{exploit}}(n) = \begin{cases} 0 & \text{if } n < 2 \\ n & \text{if } n \geq 2 \end{cases} \tag{6}$$

To gain intuition, suppose δ is so low that the scope of the relationship may at best increase by one domain upon identifying a suitable project. The players initiate cooperation by exploring projects in one domain, in part driven by the anticipation of future expansions in scope. Upon identifying a suitable project, the players choose to delay its exploitation to save on the cost incurred by player 2. Instead, they explore projects within the remaining two domains, planning to exploit the first project only after discovering one more suitable project. Should the players identify two additional suitable projects within the same period, they switch to exploit all three projects. If instead they find only one suitable project, they stop exploration and permanently exploit the two projects they have identified.

This project selection rule is the only one enabling an increase in scope after identifying a first suitable project. It is feasible due to the relatively short expected duration before exploiting the project put on hold, since the players are exploring projects in two domains simultaneously, and the motivating prospect of possibly identifying two additional projects. As noted, if the players identify only one additional suitable project, they halt exploration in the third domain. Intuitively, the cost of delaying the exploitation of a project to explore projects in the third domain is relatively large. This is because the expected duration before finding another suitable project is high, given that the players are exploring projects in only one domain. Similarly, the players cannot simultaneously exploit the two suitable projects they have already identified and explore projects in the third domain, due to the absence of any further scope expansion beyond this third domain.

To conclude, Figure 1 reports the optimal project selection rule as a function of q and v when $m = 3$, $\delta = 1/9$, and $c = 1$. In the White region, the optimal project selection rule is empty. In the Blue region, the optimal project selection rule is described by Equation (10), signifying a gradual approach with a stochastically maximal relationship scope. In the Gray region, all optimal project selection rules are gradual and result in a deterministically maximal scope, while in the Green region, the project selection rule of the symmetric-benefits benchmark is optimal. Higher values of q or v lead to project selection rules that improve joint surplus.

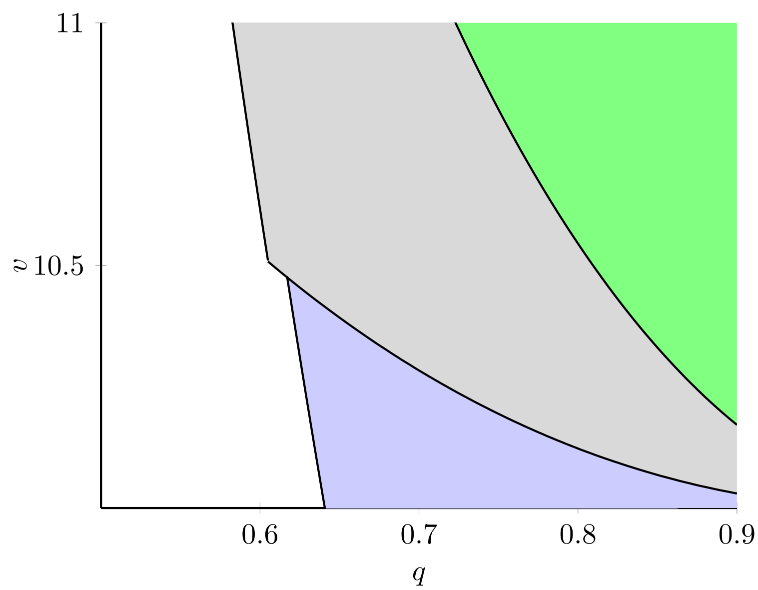

Figure 1: Optimal Project Selection Rules

We assume $m = 3$, $\delta = 1/9$, and $c = 1$. The figure starts at $v = 10$, the lowest value of v that satisfies Assumption 1. In the White region, the empty project selection rule is optimal. In the Blue region, the gradual project selection rule in which scope is stochastically maximal is optimal. In the Gray region, the optimal project selection rule is gradual and such that scope is deterministically maximal. In the Green region, the project selection rule of the symmetric-benefits case is optimal.

In this subsection, we analyzed an environment with asymmetric benefits which requires one player to be incentivized for both exploration and exploitation. In this environment, unless the discount factor is high enough to enable the implementation of the symmetric-benefits project selection rule, the scope of the players' relationship may (i) be limited in the initial phases of collaboration, (ii) gradually expand over time and (iii) fail to reach its full potential. Furthermore, when there is uncertainty about whether their collaborative relationship will attain its maximum potential scope, the players seek to expand the scope by delaying the exploitation of suitable projects.

4.2 The Dynamics of Project Exploitation Choices

In the previous subsection, we assumed that one player required incentives to both explore and exploit projects. This approach provided insights into the dynamics of the players' relationship scope. In this subsection, we assume that both players are motivated to engage in project exploration while maintaining the premise that one player requires incentives to exploit projects.[14] Additionally, we extend our analysis beyond the binary support assumption for project values, allowing us to delve into the dynamics of the players' choices about which projects to exploit.

Specifically, for any project p, we assume that $(v_{p,1} = s_p, v_{p,2} = 0)$ occurs with a probability of 1 / 2 and that $(v_{p,1} = 0, v_{p,2} = s_p)$ occurs with a probability of 1 / 2, independently across projects.[15] In addition, the distribution of s_p is subject to the following restrictions: (i) $\infty > \mathbb{E}(s_p) > 2c$, implying that exploring a project selected at random is optimal and an equilibrium of the stage game, and (ii) $\text{supp}(s_p)$ is convex to allow for a meaningful trade-off between exploration and exploitation. In this environment, scope is maximal from the beginning under any optimal relational contract, and characterizing the players' project selection rule amounts to analyzing their choices regarding which projects to exploit in all domains.

By Proposition 1, the optimal relational contract in any period t depends only on the values of the most valuable projects discovered in each of the m domains of cooperation, denoted as $\hat{s}_1, \ldots, \hat{s}_m$. Further, recall the definition of $\tilde{s}$ as established in Equation (1), which corresponds to the value a single project must generate for its

[14]In Section 5, we consider the scenario where projects benefiting both players coexist with those benefiting only a single player.

[15]This approach assumes an extreme form of asymmetric benefits. However, for our results to hold qualitatively, it suffices that each project yields a net benefit to just one player.

exploitation to be sustained on its own in equilibrium. We now provide the condition under which the players are able to follow the project selection rule of the symmetric-benefits benchmark, specifically, to exploit a project if its value is at least s^0.

Lemma 1 (**Necessary Condition for Permanent Exploitation**)
Upon finding projects with values $\hat{s}_1, \ldots, \hat{s}_m$, the players permanently follow the project selection rule of the symmetric-benefits benchmark if and only if:

$$h(\hat{s}_1, \ldots, \hat{s}_m) := \frac{1}{m} \sum_{j=1}^{m} \max\{\hat{s}_j, s^0\} \geq \tilde{s}. \tag{7}$$

When condition (7) holds, the players can pool their relational incentives to follow the project selection rule of the symmetric-benefits benchmark. This condition states that, when considering the average across all domains, the maximum value between the best project found thus far in each domain and the threshold s^0 must exceed the threshold $\tilde{s}$. We note that the function $h(\hat{s}_1, \ldots, \hat{s}_m)$ does not correspond to the arithmetic mean of the values $\hat{s}_1, \ldots, \hat{s}_m$ for two reasons: (i) the players will choose to explore rather than exploit a project with value lower than s^0 and (ii) exploration is valuable to the players and thus contributes to their continuation value.

When the condition in Lemma 1 holds and $s_j \geq s^0$ for all j, then the project selection rule from the symmetric-benefits benchmark dictates permanent exploitation, and since Equation (7) holds, the players choose "permanent exploitation". In the proposition below we show that these two conditions are, in fact, necessary and sufficient conditions. We also note that permanent exploitation occurs in finite time since, under Assumption 1, there exist project values exceeding $\tilde{s}$.

Proposition 6 (**Permanent Exploitation**)
In any optimal relational contract, the players permanently exploit projects with values $\hat{s}_1, \ldots, \hat{s}_m$ if and only if $\hat{s}_j \geq s^0$ for all $j \in \{1, \ldots, m\}$ and the average of $\hat{s}_1, \ldots, \hat{s}_m$ exceeds $\tilde{s}$.

Proof of Proposition 6.

$$\{\hat{s}_j \geq s^0\, \forall j \text{ and } \frac{\sum_j \hat{s}_j}{m} \geq \tilde{s}\} \iff \{\hat{s}_j \geq s^0\, \forall j \text{ and } \frac{\sum_j \max\{\hat{s}_j, s^0\}}{m} \geq \tilde{s}\} \qquad \square$$

The conditions listed in Proposition 6 fully characterize the players' optimal project selection rule when $m = 1$. When $m = 1$, if the players exploit project p

in period t, then $\mu_t = \mu_{t+1}$, since the players have not acquired any information during period t. It follows that the players also exploit project p in period $t+1$ and in all subsequent periods, and exploitation is thus permanent.

Corollary 1 (**Single Domain Project Selection Rule**)
When $m = 1$, in any optimal relational contract, there exists a threshold $s^*(\delta) = \max\{\tilde{s}(\delta), s^0(\delta)\}$ such that the players explore projects until they find a project p with an associated value $s_p \geq s^*$. Once they find such a project, the players exploit it in all subsequent periods.

Figure 2 provides an illustration by plotting the thresholds $\tilde{s}$, s^*, and s^0 as functions of δ, when $c = 1$ and $s_p \sim \text{Exp}(1/3)$. This distribution satisfies our assumptions, since a randomly selected project has an expected value equal to 3, which, in turn, guarantees (i) that exploration is an equilibrium of the stage game and (ii) that, for any δ, there always exists a project that can be exploited in equilibrium.

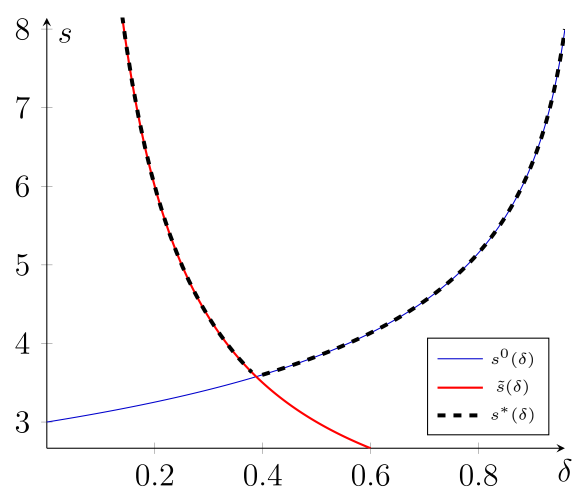

Figure 2: Comparison of Project Exploitation under Symmetric and Asymmetric Benefits

The figure plots the $\tilde{s}$, s^*, and s^0 thresholds when $s_p \sim \text{Exp}(1\,/\,3)$ and $c = 1$. The threshold $\tilde{s}$ is the minimum value of s such that cooperation in project exploitation is sustainable. The threshold s^0 is the minimum value of s such that the players switch from exploration to exploitation in the benchmark with symmetric benefits. The threshold s^* is the minimum value of s such that the players switch from exploration to exploitation. The closed-form solutions are $s^0 = 3W\Big(\delta\,/\,\big(e(1-\delta)\big)\Big) + 3$, where $W(\cdot)$ is the Lambert W function, and $\tilde{s} = 2 + (1-\delta)\,/\,\delta$. As shown in the figure, s^* is the point-wise maximum of $\tilde{s}(\delta)$ and $s^0(\delta)$.

By Proposition 2, $s^0(\delta)$ is monotonically increasing in δ and approaches infinity as $\delta \to 1$. This positive relationship arises because, in the presence of symmetric benefits, agents become more selective as their patience increases. However, from (1) it follows that $\tilde{s}(\delta)$ is monotonically decreasing in δ. This negative relationship occurs because, in the presence of asymmetric benefits, the players can cooperate in exploiting a wider range of projects as their patience grows. Therefore, a value δ^* exists where, if $\delta < \delta^*$, s^0 is lower than s^*, and the opposite is true for $\delta \geq \delta^*$. In words, for low values of δ, the players must be more selective with asymmetric compared to symmetric benefits, whereas when δ is high, they are equally selective.

We have provided closed-form conditions under which the players choose the project selection rule of the symmetric-benefits benchmark. Additionally, we have derived the conditions that dictate when projects are selected for permanent exploitation. To delve into the equilibrium dynamics that arise before the players identify a set of m projects suitable for permanent exploitation, we now suppose $m = 2$ and focus on a specific parametric example.

Exponential Distribution Example When $m = 2$. Suppose $c = 1$, $\delta = 1 \,/\, 3$, and $s_p \sim (3 - 1 \,/\, \lambda) + \text{Exp}(\lambda)$, such that $\mathbb{E}\left(s_p\right) = 3$. Further, denote by $\mathcal{C}(\hat{s}_1, \hat{s}_2)$ the continuation value of the players' relationship as a function of $\hat{s}_1$ and $\hat{s}_2$. We first note that $1 \leq \mathcal{C}(\hat{s}_1, \hat{s}_2)$ for all combinations of $\hat{s}_1$ and $\hat{s}_2$. This inequality holds because the players can always choose to explore two new projects in every period, generating a payoff of $3 - 2$ per project and thus a continuation value of $\frac{\delta}{1-\delta}2$, which reduces to 1 when $\delta = 1 \,/\, 3$. Within this example, (i) we utilize Proposition 6 to characterize the players' project selection rule when $h(\hat{s}_1, \hat{s}_2) \geq \tilde{s}$ and (ii) we derive results for the optimal project selection rule when $h(\hat{s}_1, \hat{s}_2) < \tilde{s}$.

Figure 3a displays the threshold s^0 as dotted black lines and the set of $\hat{s}_1$ and $\hat{s}_2$ values that satisfy $h(\hat{s}_1, \hat{s}_2) = \tilde{s}$ in red. This set is the solution to Equation (7). In Figure 3b, we indicate which projects are chosen for exploitation based on the respective values of $\hat{s}_1$ and $\hat{s}_2$. Each region is denoted by the set of domains for which the highest-valued project in that domain is exploited. First, it follows from Figure 3a that both projects are chosen for exploitation when $h(\hat{s}_1, \hat{s}_2) \geq \tilde{s}$ and $\hat{s}_1, \hat{s}_2 \geq s_0$. Outside of this region, we must address two questions: (i) will there be a project selected for exploitation, and (ii) if so, which one? Evidently, the answer to the second question is the best project. Therefore, in Figure 3b, the selection between $\hat{s}_1$

and $\hat{s}_2$ hinges on which side of the 45 degree line the project values fall. Further, one can prove that there exists a threshold, s', on the value of the best of the two projects such that, below this threshold, the players choose to explore two new projects rather than exploiting the best of the two projects. Conversely, above this threshold, the players choose to exploit the best of the two projects and explore a new one. Moreover, the threshold s' is independent of the value of the worse of the two projects, since this project will never be exploited. Finally, Figure 3b also presents a sample path illustrating the evolution of realized project values over time, depicted in blue. In the phase where the players are exploring two projects simultaneously, both $\hat{s}_1$ and $\hat{s}_2$ weakly increase over time. In the phase where the players exploit a project in domain j, $\hat{s}_j$ remains constant, while $\hat{s}_{-j}$ weakly increases over time. Finally, in the phase of the relationship where the players exploit both projects, $\hat{s}_1, \hat{s}_2$ stay constant because exploitation is permanent. Arrows are used to signify changes in project values when a more valuable project is identified, while self-loops indicate situations where more valuable projects are either not discovered or not pursued.

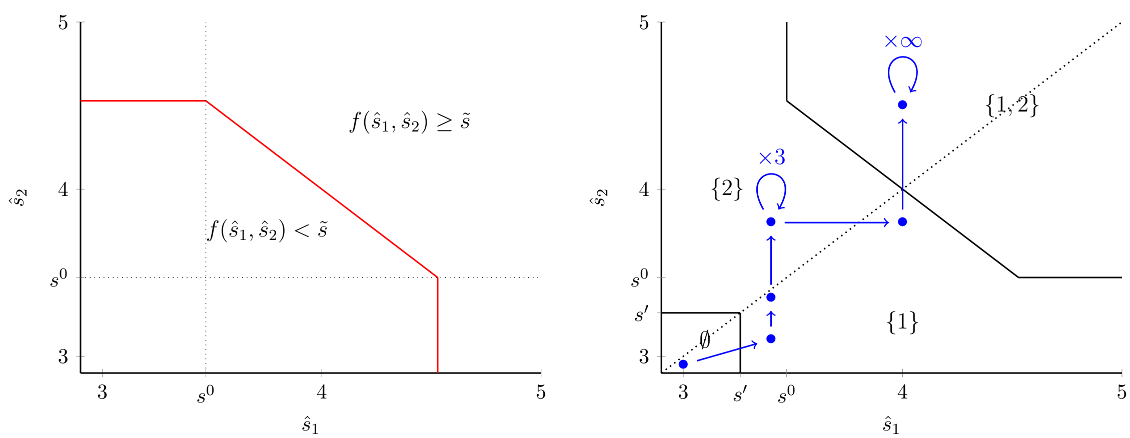

(a) Feasible Region for First-Best Project Exploitation

(b) Project Exploitation and Sample Path

Figure 3: Optimal Multi-Project Selection Dynamics

In the figure, we assume $c = 1$, $m = 2$, $\delta = 1/3$, and $s_p \sim \text{Exp}(1/3)$. The values $\hat{s}_1$ and $\hat{s}_2$ represent the values of the best projects discovered by the players to date in cooperative domains 1 and 2, respectively. The left figure plots (i) the threshold s^0 such that the players switch from exploration to exploitation in the benchmark case with symmetric benefits and (ii) the set of $\hat{s}_1$ and $\hat{s}_2$ values such that $h(\hat{s}_1, \hat{s}_2) = \tilde{s}$ in red. The right figure plots in Black the project selection behavior of the players under the optimal relational contract. Each region is denoted by the set of domains for which the highest-valued project in that domain is exploited. For example, in the region $\{1, 2\}$, the best project for each domain is exploited. In region $\{\emptyset\}$, no project is chosen for exploitation. In Blue, we plot one realization of a sample path.

The sample path of realized project values depicted in Figure 3b implies novel dynamics along the equilibrium path, characterized by several instances of temporary project exploitation, including of projects with values below s^0.[16] Building on this observation, in what follows we assume that $s_p \sim (3 - 1 / \lambda) + \text{Exp}(\lambda)$ without restricting δ to $1/3$ or c to 1 . In this environment, we provide an analytical proof that, beyond the departures already highlighted in Proposition 6, additional discrepancies from the symmetric-benefits benchmark can arise in equilibrium.

Proposition 7 (**Further Departures From Symmetric Benefits Rule**)
Suppose $s_p \sim (3 - 1 / \lambda) + \text{Exp}(\lambda)$ and $m = 2$. For an open set of parameter values, the following behaviors occur with positive probability on path:

1. *The players exploit a project in period t but not in some period $t' > t$.*

2. *The players exploit a project with value $s_p < s^0$.*

3. *The players choose not to exploit a project in period t, but choose to exploit it in some later period $t' > t$.*

The behaviors described in the proposition do not occur with positive probability across all parameter values, a conclusion that follows when noting that, as δ converges to 1, players follow the project selection rule of the symmetric-benefits benchmark. Furthermore, even within the parameter range where these behaviors occur with a positive probability, their occurrence is not guaranteed: for instance, there is always a non-zero probability that the players immediately identify two projects worthy of permanent exploitation.

As argued above, the first statement follows from the second statement combined with the result from Proposition 6, whereby the players only permanently exploit projects whose values exceed s^0. The intuition behind the second statement can be seen by comparing the players' exploration incentives in the presence of symmetric versus asymmetric benefits, and by supposing that s_p is distributed such that the continuation value of the players' relationship exceeds $2c$ with an arbitrarily small probability. In this case, with a probability approaching 1, the players are able to exploit at most one project, implying the exploration of at least one other project.

[16] In this scenario, the players never return to previously abandoned projects. However, when $m > 2$, players may opt to exploit a project for several periods, subsequently abandon it, and later revert to its exploitation.

Consequently, the benefit of exploration stems from the chance to find a more valuable project for exploitation in the next period. However, since the players are already exploring one project, the gain from undertaking a second exploration might be limited: even if two superior projects are identified, they can only exploit one. On the other hand, with projects generating symmetric benefits, the players, upon discovering these two projects, can exploit both, providing them with greater exploration incentives. This higher benefit of exploration, in turn, prompts the players to set a higher exploitation threshold when benefits are symmetric.

Regarding the final statement, consider a value of δ sufficiently small such that the players are unable to cooperate in exploiting projects achieving values slightly exceeding s^0. If in period 0 the players do find two projects with associated values only slightly higher than s^0, the players are compelled to explore two new projects during the next period. However, if one of these new projects happens to achieve a high value, the continuation value of the players' relationship may exceed $2c$ and the players may wish to return to one of the two period 0 projects.

In summary, this subsection has shown that in scenarios characterized by asymmetric project benefits, when both players are motivated to explore projects but one player requires incentives to engage in project exploitation, the players are likely to explore projects for a considerable amount of time before identifying suitable projects for permanent exploitation. Furthermore, players engage in incentive pooling across different cooperation domains to facilitate collaboration, resulting in significant path dependence in the selection of projects for exploitation. Throughout the collaboration, players may interrupt the exploitation of a project in one domain upon discovering a superior project in another domain, and they may revisit previously explored projects in some domains when valuable projects are found in other domains.

5 Extensions

In this section, we provide additional results regarding the role played by the dimensionality m of the players' relationship. Next, we explore two extensions of the main model, incorporating projects with both symmetric and asymmetric benefits.

5.1 On the Benefits of Scope

Denote by $\tilde{\pi}(m) := \pi(m) \,/\, m$ the average joint surplus of the relationship per domain of the relationship. Similarly, denote by $\delta^*(m)$ the minimum discount factor for which the optimal relational contract is non-empty. The following two inequalities hold: $\tilde{\pi}(m \cdot k) \geq \tilde{\pi}(m)$ and $\delta^*(m \cdot k) \leq \delta^*(m)$, with $k \geq 1$. The intuition behind these weak inequalities is that the players can always engage in k independent relationships, each replicating the optimal relational contract with m domains.

In the environment considered in Section 4.1, where exploration is not an equilibrium of the stage game, $\delta^*(m \cdot k) < \delta^*(m)$. To see why this inequality holds strictly, recall that, when $\delta = \delta^*(m)$, the players' optimal relational contract is necessarily gradual. The players could adopt k separate gradual relational contracts. However, doing so would be inefficient because it would condition further project exploration exclusively on the number of projects suitable for exploitation found within an (inefficiently segmented) separate relational contract. By the same argument, $\tilde{\pi}(m \cdot k) > \tilde{\pi}(m)$ whenever the optimal relational contract exhibits gradualism.[17]

In the environment considered in Section 4.2, where exploration can always be sustained in equilibrium, $\delta^*(m) = 0$. However, one can show that $\tilde{\pi}(m \cdot k) \geq \tilde{\pi}(m)$ holds as a strict inequality in the range of parameter values such that the players are unable to replicate the project selection rule of the symmetric-benefits benchmark. In these instances, pooling relational incentives across the previously k independent relationships is valuable to the players.

5.2 Symmetric Projects as Stepping Stones in Relationships

We can also explore the dynamics of the scope of the players' relationship when projects with both symmetric and asymmetric benefits are available. Suppose that the distribution of project benefits is as in Section 4.1, namely $v_{p,2} = 0 \, \forall p \in \mathcal{P}$. However, the players now also have access to m projects (each associated with a distinct domain) with guaranteed benefits of v_c for both players. We assume that $2c < 2v_c < \mathbb{E}(s_p)$ to ensure that these projects are profitable, but less so (in expectation) compared to

[17]While $\pi(m)$ is monotonically increasing in m, $\tilde{\pi}(m)$ may not be. To see this, suppose that s_p belongs to a three-point support (low, medium, and high). Suppose further that a high-valued project and a single medium-valued project can be jointly exploited by the players, but that a high-valued project and two medium-valued projects cannot, then $\tilde{\pi}(m)$ would depend on the parity of m and monotonicity would break.

those chosen from the set $\mathcal{P}$. As the projects with symmetric benefits can always be selected, the presence of projects with symmetric benefits increases the value of the players' relationship for two related reasons: (i) the players' relationship is maximal from the very beginning and (ii) the presence of valuable projects with symmetric benefits allows the players to begin exploring and subsequently exploiting the more profitable asymmetric projects earlier. In this sense, the symmetric projects act as stepping stones in the building of the players' relationship.

5.3 Favoring Symmetric over Asymmetric Projects

In the Online Appendix, we modify the setting of Section 4.2 to allow for projects with both symmetric and asymmetric benefits across the players. We show that, when $m = 1$, players exhibit less selectivity for projects with symmetric benefits compared to asymmetric ones. Additionally, the existence of asymmetric projects diminishes selectivity for symmetric projects compared to when only symmetric projects are available. Two distinct exploration/exploitation thresholds exist, s_s^* and s_a^*, where the first threshold applies to projects with symmetric benefits, while the second one applies to those with asymmetric benefits. The key finding that emerges from this analysis is that: $s_a^* \geq s^0 \geq s_s^*$. The intuition behind $s_a^* \geq s^0$ is the same intuition as before: a project that benefits just one player must be valuable enough to enable cooperation in exploitation. The intuition behind $s^0 \geq s_s^*$ is as follows. If the players identify a project with value $s_p \in (s^0, s_a^*)$, they can exploit it only if it yields symmetric benefits. As a result, the overall value of exploration is lower for the players, leading to lower exploration/exploitation thresholds for projects with symmetric benefits. These results bear resemblance to Acharya and Ortner (2022), where two players involved in collective search prioritize projects that benefit both players.

6 Applications

In this section, we integrate our theoretical analysis with case studies and empirical research focused on collaborative dynamics. First, we consider inter-firm collaborations, with an emphasis on buyer-supplier relationships. Next we examine how our framework can also be applied to understand intra-firm collaborative dynamics, such as interactions between managers and employees, and connect our findings with the literature on persistent performance differences across firms.

6.1 Between Firms

We have shown that when asymmetric benefits are present among collaborating parties and the value of the relationship is initially low, parties may find it necessary to: (i) expand the scope of their collaboration only gradually, leveraging early successes as stepping stones to broaden their partnership, as discussed in Section 4.1; and (ii) to engage in prolonged experimentation to identify projects with substantial value, so that no party will find it in their interest to withdraw their cooperation in the long run, as detailed in Section 4.2.

These predictions and their underlying logic align closely with the applied literature focusing on buyer-supplier dynamics. In buyer-supplier relationships, suppliers often undertake non-contractible investments that tend to benefit the buyers. Toyota's relationship with its suppliers is a well-documented example of a gradual and experimental approach, where Toyota encouraged each of its suppliers to incrementally adopt various practices from the Toyota Production System—a strategy that initially yielded benefits primarily for Toyota through improvements in quality and efficiency (Dyer and Nobeoka, 2000).

Moving beyond the Toyota example, Vanpoucke et al. (2014) identify commonalities across case studies through a survey involving more than one hundred buyer–supplier relationships. They find that these relationships often exhibit gradualism and lengthy experimentation. For example, analyzing a collaboration involving the development of soy bean products, Vanpoucke et al. (2014) observe that "the buyer decided to purchase from this supplier. This was the start of the exploration stage. ... it took about 10 years before the two parties started up a first integration initiative and entered the expansion stage. It then took a long time for both partners to get to know each other, build enough trust and see the benefits of learning from each other's expertise and further building up the relationship." Consistent with our analysis, Vanpoucke et al. (2014) emphasize the strong path dependence that marks relationship dynamics, noting that "events rather than time define the development stage of the relationship." In the previous example regarding soy bean products, and across the other cases analyzed in Vanpoucke et al. (2014), the impetus for additional joint collaborations came from successes in previous domains of cooperation.

In this applied literature, the need for collaborating parties to build the credibility necessary to sustain their relationship is often cited as a key explanation for gradualism and lengthy experimentation. In a McKinsey report about buyer-supplier dynamics, Gutierrez et al. (2020) note that "Building trust takes time [...]. Often this means starting small, with simple collaboration efforts that deliver results quickly, building momentum." Similarly, Dwyer et al. (1987) note that "the critical distinction [between the exploration and expansion phases in buyer-supplier relationships] is that the rudiments of trust and joint satisfactions established in the exploration stage now lead to increased risk taking within the dyad. Consequently, the range and depth of mutual dependence increase."

Further, our analysis sheds light on patterns documented by an empirical literature in economics on the strong persistence of buyer-supplier relationships and the tendency for relationship value to increase over time (e.g., Macchiavello and Morjaria, 2015; Bernard and Moxnes, 2018; Monarch, 2022; Monarch and Schmidt-Eisenlohr, 2023). The persistence of relationships is surprising according to existing models that suggest suppliers would gain significantly in productivity by changing partners frequently, and can only be explained if significant exogenous costs associated with switching or searching for new partners are assumed. Our analysis suggests that buyers may be reluctant to switch suppliers, even when alternatives seem better, precisely because relationship value grows over time. To be certain, growth in relationship value can occur for a number of reasons. Learning dynamics, such as instances where suppliers gain confidence in their partners' ability to fulfill orders or sincere intention to collaborate, can contribute to the enhancement of relationship value, lead to larger transaction volumes, and foster the persistence of the relationship (c.f., Rauch and Watson, 2003). Our framework shifts focus towards a complementary mechanism for the growth of relationship value: the gradual expansion of collaboration scope between buyers and suppliers across domains of cooperation. This expansion in scope offers an additional explanation for the sustained increase in the value of buyer-supplier relationships, maintaining its significance even once both parties have established confidence in each other's capabilities and intentions.

6.2 Within Firms

While much of our focus has been on interactions between firms, our model serves as a valuable lens for examining employer-employee dynamics. In this context, one can conceptualize one party in our model as the employer and the other as the employee, where, for instance, benefits consistently accrue to the employer, as assumed in Section 4.1. Furthermore, the different domains of collaboration can be seen as various responsibilities assigned to the employee, with every project representing a potential managerial approach for collaborative task achievement.

With this interpretation in mind, our work also contributes to the literature on firm performance and, specifically, persistent performance differences among seemingly similar enterprises (see Syverson, 2011; Gibbons and Henderson, 2013, and references therein). Numerous empirical studies have documented enduring disparities in performance across a range of industries and countries, and these gaps have proven surprisingly robust against plausible explanations such as market competition or local geographical and demand conditions. According to Gibbons and Henderson (2013), and the body of evidence they review, variations in managerial practices are key in creating productivity disparities across firms. We adapt for our purposes their categorization of explanations: (i) managers might either be unaware of their poor performance, or, even if aware, believe that the best practices from other firms are not suitable for their context; (ii) managers are aware of their poor performance and are able to seek superior managerial practices suitable to their context, but opt not to; and (iii) managers are "striving mightily" to adopt superior practices but face hurdles during the implementation phase. The first explanation underscores information barriers, prompting questions about why such information does not diffuse more readily (c.f. Bloom et al., 2013; Atkin et al., 2017). The second explanation is consistent with the framework developed by Chassang (2010) and discussed in Section 1.1, in which players are informed about the existence of more efficient practices but choose not to pursue them. This explanation is also consistent with our analysis in Section 4.1, showing that not all collaborations are guaranteed to reach their maximum potential despite identical initial conditions (Proposition 5).

Our model also offers insight into the third explanation presented by Gibbons and Henderson (2013). To illustrate, consider a scenario with two ex ante identical organizations, each consisting of one manager and one employee collaborating over multiple domains (e.g., sales and customer support), and a low discount factor δ.

Both organizations are in the process of improving their managerial practices in each domain through trial-and-error, and solutions are organization-specific. The low discount factor compels the parties in both organizations to initially focus their efforts on identifying a superior managerial practice within a single domain of their collaboration. They do so despite being aware that superior practices exist in other domains as well. If one organization happens to discover a successful managerial practice in the initial domain of focus early on, while the other organization does not, the paths of the two organizations will diverge. The first organization, having found success early, will be able to move on and seek superior managerial practices within additional domains. Consequently, the first organization has the potential to persistently maintain superior performance compared to the second organization. This performance gap is expected to widen over extended periods due to the compounding impact of each successful managerial practice discovered and implemented by the first organization. As they continue to find and adopt superior practices across multiple domains, their overall performance will continue to improve, leaving the less fortunate organization further behind. The second organization is "striving mightily" (and may eventually succeed) to match the performance of the first organization, as it is still attempting to achieve success in the first domain. However, identifying superior practices is a time-intensive endeavor, and the second organization will not be able to increase its scope until finding such a success.

7 Concluding Remarks

This paper has presented a framework for examining the dynamics of multi-project collaborations, particularly when benefits are distributed asymmetrically among the parties. The model generates three key insights. First, in situations where relationship value is initially low, the parties cannot immediately realize the full potential scope of their relationship. To build the credibility needed for sustaining cooperation across multiple projects, the parties start by cooperating on a select few projects. Successes achieved in these initial projects serve as a foundation for expanding their collaborative efforts to additional domains. Second, because credibility is intricately tied to the value of the players' relationship, collaborating parties invest a substantial amount of time in seeking projects that are sufficiently valuable to enable cooperation. Third, parties combine their relational incentives across projects, resulting in inter-dependencies and significant path dependence. In this context,"relational inter-dependencies" between projects can lead to seemingly erratic behaviors, such as prolonged cooperation on projects that are ultimately discontinued, or the revival of previously abandoned projects, all driven by the time required for the parties to build credibility in their relationship.

References

Abreu, Dilip, "Extremal equilibria of oligopolistic supergames," *Journal of Economic Theory*, 1986, *39* (1), 191–225.

Acharya, Avidit and Juan Ortner, "Paths to the Frontier," *American Economic Journal: Microeconomics*, 2022, *14* (1), 39–69.

Anesi, Vincent and T. Renee Bowen, "Policy Experimentation in Committees: A Case against Veto Rights under Redistributive Constraints," *American Economic Journal: Microeconomics*, August 2021, *13* (3), 124–62.

Atkin, David, Azam Chaudhry, Shamyla Chaudry, Amit K. Khandelwal, and Eric Verhoogen, "Organizational Barriers to Technology Adoption: Evidence from Soccer-Ball Producers in Pakistan," *The Quarterly Journal of Economics*, 03 2017, *132* (3), 1101–1164.

Baker, George, Robert Gibbons, and Kevin J. Murphy, "Subjective Performance Measures in Optimal Incentive Contracts," *The Quarterly Journal of Economics*, 1994, *109* (4), 1125–1156.

_ , _ , and _ , "Relational Contracts and the Theory of the Firm," *The Quarterly Journal of Economics*, 2002, *117* (1), 39–84.

Benavides, Luis, Verda De Eskinazis, and Daniel Swan, "Six Steps to Successful Supply Chain Collaboration," *McKinsey*, July 2012.

Bergemann, Dirk and Juuso Välimäki, "Bandit Problems," in Steven Durlauf and Larry Blume, eds., *The New Palgrave Dictionary of Economics*, 2nd ed., Macmillan Press, 2008.

Bernard, Andrew B. and Andreas Moxnes, “Networks and Trade,” *Annual Review of Economics*, 2018, *10* (1), 65–85.

Bernheim, B Douglas and Michael D Whinston, “Multimarket contact and collusive behavior,” *The RAND Journal of Economics*, 1990, pp. 1–26.

Bloom, Nicholas, Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts, “ Does Management Matter? Evidence from India,” *The Quarterly Journal of Economics*, 11 2013, *128* (1), 1–51.

Bolton, Patrick and Christopher Harris, “Strategic experimentation,” *Econometrica*, 1999, *67* (2), 349–374.

Bonatti, Alessandro and Johannes Hörner, “Collaborating,” *American Economic Review*, 2011, *101* (2), 632–663.

Bull, Clive, “The Existence of Self-Enforcing Implicit Contracts,” *The Quarterly Journal of Economics*, 1987, *102* (1), 147–159.

Chan, Jimmy, Alessandro Lizzeri, Wing Suen, and Leeat Yariv, “Deliberating collective decisions,” *The Review of Economic Studies*, 2018, *85* (2), 929–963.

Chassang, Sylvain, “Building routines: Learning, cooperation, and the dynamics of incomplete relational contracts,” *American Economic Review*, 2010, *100* (1), 448–65.

Dwyer, F. Robert, Paul H. Schurr, and Sejo Oh, “Developing Buyer-Seller Relationships,” *Journal of Marketing*, 1987, *51* (2), 11–27.

Dyer, Jeffrey H and Kentaro Nobeoka, “Creating and managing a high-performance knowledge-sharing network: The Toyota case,” *Strategic Management Journal*, 2000, *21* (3), 345–367.

Ghosh, Parikshit and Debraj Ray, “Cooperation in Community Interaction without Information Flows,” *The Review of Economic Studies*, 1996, *63* (3), 491–519.

Gibbons, Roberts and Rebecca Henderson, “What Do Managers Do? Exploring Persistent Performance Differences Among Seemingly Similar Enterprises,” *Handbook of Organizational Economics*, 2013, pp. 342–372.

Gutierrez, Agustin, Ashish Kothari, Carolina Mazuera, Tobias Schoenherr, and McKinsey Company, "Taking supplier collaboration to the next level," *McKinsey Company*, 2020.

Halac, Marina, "Relationship building: conflict and project choice over time," *The Journal of Law, Economics, & Organization*, 2014, *30* (4), 683–708.

_ , "Investing in a Relationship," *The RAND Journal of Economics*, 2015, *46* (1), 165–185.

Hörner, Johannes, Nicolas Klein, and Sven Rady, "Overcoming Free-Riding in Bandit Games," *The Review of Economic Studies*, 11 2022, *89* (4), 1948 – 1992.

Ide, Enrique, "Dual Moral Hazard and The Tyranny of Success," American Economic Journal: Microeconomics 2024. Forthcoming.

Keller, Godfrey, Sven Rady, and Martin Cripps, "Strategic Experimentation with Exponential Bandits," *Econometrica*, 2005, *73* (1), 39–68.

Kranton, Rachel E., "The Formation of Cooperative Relationships," *Journal of Law, Economics, Organization*, 1996, *12* (1), 214–233.

Levin, Jonathan, "Multilateral contracting and the employment relationship," *The Quarterly Journal of Economics*, 2002, *117* (3), 1075–1103.

_ , "Relational Incentive Contracts," *American Economic Review*, June 2003, *93* (3), 835–857.

Macchiavello, Rocco and Ameet Morjaria, "The value of relationships: evidence from a supply shock to Kenyan rose exports," *American Economic Review*, 2015, *105* (9), 2911–2945.

Macleod, W. Bentley and James Malcomson, "Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment," *Econometrica*, 1989, *57* (2), 447–80.

Monarch, Ryan, "It's Not You, It's Me: Prices, Quality, and Switching in U.S.-China Trade Relationships," *The Review of Economics and Statistics*, 09 2022, *104* (5), 909–928.

— **and Tim Schmidt-Eisenlohr**, "Longevity and the value of trade relationships," *Journal of International Economics*, 2023, *145*, 103842.

Rauch, James E. and Joel Watson, "Starting small in an unfamiliar environment," *International Journal of Industrial Organization*, 2003, *21* (7), 1021–1042.

Reshidi, Pëllumb, Alessandro Lizzeri, Leeat Yariv, Jimmy H Chan, and Wing Suen, "Individual and collective information acquisition: An experimental study," Technical Report, National Bureau of Economic Research 2021.

Strulovici, Bruno, "Learning while voting: Determinants of collective experimentation," *Econometrica*, 2010, *78* (3), 933–971.

Syverson, Chad, "What Determines Productivity?," *Journal of Economic Literature*, June 2011, *49* (2), 326–65.

Tirole, Jean, *A Theory of Corporate Finance*, Princeton University Press, 2006.

Urgun, Can, "Restless Contracting," 2021. Working Paper.

Vanpoucke, Evelyne, Ann Vereecke, and Kenneth K. Boyer, "Triggers and patterns of integration initiatives in successful buyer–supplier relationships," *Journal of Operations Management*, 2014, *32* (1), 15–33.

Venables, Sarah, "Essays on organizational economics and the heterogeneity of firm performance." Ph.d. dissertation, Massachusetts Institute of Technology, Cambridge, MA, USA 2013.

Watson, Joel, "Starting small and renegotiation," *Journal of Economic Theory*, 1999, *85* (1), 52–90.

— , "Starting Small and Commitment," *Games and Economic Behavior*, 2002, *38* (1), 176–199.

Weitzman, Martin, "Optimal Search for the Best Alternative," *Econometrica*, 1979, *47* (3), 641–654.

A Appendix

Proof of Proposition 2.

Recall from the text that one can ignore Equation (2) when analyzing the equilibrium that maximizes the joint surplus of the players. As such, since there are no inter-depenencies across domains, the players will treat each domain identically and symmetrically.

Statement 1: All projects in $\mathcal{P}_j$ are ex ante identical. Hence, if the players ever find it optimal to explore a project, then in every period the players will select a project from domain j. This implies that for each domain the players either choose no projects in all periods or a project in every period.

Statement 2: Given that there exists an infinite number of ex-ante identical projects, the optimal relational contract conditions only on the project with the highest value amongst all previously explored projects, whose value we denote $\hat{s}$. In particular, one can write the Bellman Equation:

$$B(\hat{s}) = \max_{\text{explore, exploit } \hat{s}} \{\mathbb{E}\left(s'\right) - 2c + \delta\mathbb{E}\left(B\left(\max\left(\hat{s}, s'\right)\right)\right), \hat{s} - 2c + \delta B\left(\hat{s}\right)\}. \tag{8}$$

The first term in the maximum operator corresponds to the players' expected surplus when exploring one more project (chosen at random, since all unexplored projects are ex ante identical) and the second term is their surplus when exploiting the best project found thus far. Next, one can show there exists a threshold s^0, wherein the players explore if $\hat{s} < s^0$ and exploit if $\hat{s} \geq s^0$. Finally, for any $\delta < 1$, one can use Blackwell's Sufficient Conditions to show that there exists a unique solution to the Bellman Equation, and hence the threshold rule dictated by s^0 is a solution. $\square$

Proof of Proposition 4. By Equation (2), under any optimal relational contract specifying $|\mathbf{P}^0| = m$, the continuation value $\mathcal{C}$ exceeds $c \cdot m$ at date 0 and can only grow over time as n increases. The players therefore follow the project selection rule of the symmetric benefits benchmark and $\mathcal{C}$ is additively separable across domains. Further, as the scope remains constant under the symmetric benefits benchmark, and the continuation value is lowest in period 0, the most binding incentive constraint is the one for period 0 provided in Equation (3). Because $\mathcal{C}(\text{exploration}) = 0$ when $\delta = 0$ and is strictly increasing in δ, there exists a threshold $\overline{\delta}$ such that (3) fails to hold when $\delta < \overline{\delta}$. Thus, if $\delta \geq \overline{\delta}$, any optimal relational contract specifies the project selection rule of the symmetric benefits benchmark and $|\mathbf{P}^0| = m$.

We now show the existence of a non-empty range of discount factors within which any optimal relational contract specifies $0 < |\mathbf{P}^0| < m$. We show this result by considering an arbitrary project selection rule with $0 < |\mathbf{P}^0| < m$. Specifically, consider (i) $f_{\text{explore}}(0) = 1$ and $f_{\text{explore}}(n) = m - n$ if $n \geq 1$ and (ii) $f_{\text{exploit}}(n) = n$. There exists a relational contract that implements this project selection rule if and only if Inequality (2) holds when $n = 0$ and $n = 1$, which are provided as Equation (4) and (5), respectively. As argued in the text, these inequalities are strictly easier to satisfy and thus the gradual project selection rule has a strictly lower associated critical discount factor $\bar{\delta} < \delta^*$.

Finally, notice that Inequality (3) is independent of m. Consequently, when the players opt for a project selection rule where $0 < |\mathbf{P}^0| < m$, scope must necessarily increase along the equilibrium path. To see this, assume, by contradiction, that the scope remains constant and equal to $|\mathbf{P}^0|$. Then, Inequality (2) at date 0 coincides with Inequality (3) (upon replacing m with $|\mathbf{P}^0|$). However, if Inequality (3) holds, the players adopt the project selection rule of the symmetric-benefits benchmark. $\square$

Proof of Proposition 5. **Statement 1:** We define n as an absorbing state if n is reached with positive probability and, upon discovering n projects suitable for exploitation, the players choose to exploit all n projects, and thus have a scope that remains at n for all subsequent periods. We begin by showing that no absorbing state can be reached with probability one.

Let's suppose, by contradiction, that $0 < n < m$ is an absorbing state reached with probability 1. No project selection rule can guarantee the players to identify n projects suitable for exploitation without previously having identified $n - 1$ such projects. Consider the case in which the players have identified $n-1$ projects suitable for exploitation. In order for the players to identify precisely one more project suitable for exploitation, they must explore projects on exactly one additional domain of cooperation. Additionally, according to Proposition 3, any optimal relational contract is such that $f(n) \geq f(n-1)$ and $f(n-1) \geq n-1$. Thus, the chain of inequalities $n - 1 \leq f(n-1) \leq f(n) = n$ must hold. This means that the project selection rule for $n - 1$ takes one of two possible forms. Suppose first that the players exploit the $n-1$ projects suitable for exploitation they have found thus far and that they explore projects in one additional domain. For this project selection rule to be implementable, by Proposition 1 the following inequality must hold:

$$nc \leq \mathcal{C}(\text{exploration}) + (n-1)\mathcal{C}(\text{exploitation}). \tag{9}$$

However, Assumption 1 implies that $\mathcal{C}(\text{exploitation}) > c$, so that Inequality (9) is relaxed a n increases. Thus, if Inequality (9) holds, upon identifying an nth project suitable for exploitation, the players would be able to exploit all n projects and explore projects in one additional domain. This scenario leads to a contradiction. Hence, when the players have identified $n-1$ projects suitable for exploitation, it must be that they explore projects in one additional domain and exploit $n-2$ projects. However, note that exploiting $n-1$ projects is always feasible. Following this project selection rule therefore implies that the players prefer the exploration of projects in one additional domain of cooperation and eventual exploitation of two projects to the immediate and permanent exploitation of a single project. Yet, once the players have identified n projects suitable for exploitation, they face the same trade-off. Given that they cease exploring projects after identifying n suitable ones, it follows that they would also stop exploring once they have identified $n-1$ such projects. Thus we know no absorbing state is reached with probability one, and further if n is an absorbing state, the players conduct two explorations when $n-1$ suitable projects have been discovered.

Given this result, it is sufficient to show that n and $n+1 < m$ cannot both be absorbing states. To see why, note that if n is an absorbing state, but $n+1$ is not, then the previous result shows that if the relationship reaches $n-1$ the players conduct two explorations, and thus reach $n+1$ with positive probability. Thus, if there are never two consecutive absorbing states, the relationship reaches maximal scope with positive probability. Let us now assume by contradiction that both n and $n+1 < m$ are absorbing. Given this assumption, at $n-1$ the players conduct two explorations and $n-2$ exploitations. As the players could always choose to conduct $n-1$ explorations, the players prefer to forego one exploitation for two explorations.

Next, note that if at $n-1$ the players could conduct $n-2$ exploitations and two explorations , the players can similarly conduct $n-1$ exploitations and two explorations at n (by a similar observation to Equation (9) being increasing in n). Further, as the players preferred to sacrifice one exploitation for two explorations previously, the players will have the same preference at n. This is a contradiction, as we previously assumed n was an absorbing state. $\square$

Statement 2: The proof for this statement was provided in the text. $\square$

Statement 3: By statement 2, we must consider $m \geq 3$. Also, by an identical reasoning to that provided in the proof of statement 2, we know that at $n = 1$ the players must conduct an exploration. Suppose $m = 3$ and consider 2 as an absorbing state of the players' relationship (as defined in the proof of statement 1). The proof of Figure 1 is included in the Online Appendix.[18] We show that the following project selection rule is the only one inducing an absorbing state of $n = 2$:

$$f_{\text{explore}}(n) = \begin{cases} n+1 & \text{if } n \leq 1 \\ 0 & \text{if } n > 1 \end{cases} \quad f_{\text{exploit}}(n) = \begin{cases} 0 & \text{if } n < 2 \\ n & \text{if } n \geq 2 \end{cases} \tag{10}$$

One can note from (10) that with probability $\frac{q}{2-q}$ $\lim_{t\to\infty} |\mathbf{P}^t| = 3$ and with complementary probability $\frac{2-2q}{2-q}$ $\lim_{t\to\infty} |\mathbf{P}^t| = 2$. $\square$ $\square$

Proof of Lemma 1. When the players have identified projects with values $\hat{s}_1, \ldots, \hat{s}_m$ at history h, the condition for the players being able to replicate the project selection rule of the symmetric-benefits benchmark in all subsequent periods is that, for all histories h' occurring after h and with associated project values $\hat{s}'_1, \ldots, \hat{s}'_m$, the players exploit $\hat{s}'_j$ if and only if $\hat{s}'_j \geq s^0$. This condition is as follows:

$$c \sum_{j=1}^{m} \mathbf{1}_{\hat{s}'_j \geq s^0} \leq \delta \left(\sum_{j=1}^{m} \mathcal{C}(\hat{s}'_j) \right) \forall (\hat{s}'_1, \ldots, \hat{s}'_m) \geq (\hat{s}_1, \ldots, \hat{s}_m), \tag{11}$$

which corresponds to Equation (2) when the players follow the symmetric-benefits benchmark where $\mathcal{C}(\hat{s}'_j)$ denotes the continuation value in domain j under the symmetric-benefits benchmark. Note that such a function is (i) constant below s^0, (ii) $\lim_{x\uparrow s^0} \mathcal{C}(x) > \lim_{x\downarrow s^0} \mathcal{C}(x)$ and (iii) increasing after s^0. Given such properties, one can note that setting $\hat{s}'_i = \max\{\hat{s}_i, s^0\}$ both minimizes the right-hand side and maximizes the left-hand side of Equation (11). Thus, an equivalent condition is:

$$m \cdot c \leq \delta \Big(\sum_{j=1}^{m} \frac{1}{1-\delta} (\max\{\hat{s}_j, s^0\} - 2) \Big), \tag{12}$$

[18]See Footnote 13 for why such characterizations are lengthy and are thus included in the Online Appendix.

which corresponds to the expression stated in the Lemma. □

Proof of Proposition 7. Suppose $s_p \sim (3-1/\lambda)+\text{Exp}(\lambda)$, which implies that $\mathbb{E}(s_p) = 3 > 2$. Moreover, the support of this distribution is convex for any λ. Hence, this distribution satisfies the assumptions made in the text. Recall that the optimal relational contract conditions only on the best project found thus far in each domain.

Statement 3 Note that there exists a sufficiently small value of δ such that the players are unable to exploit a project worth $s^0+\epsilon$. Consider such a δ. With positive probability, in period 0 the players identify two projects with values belonging to an arbitrarily small range around $s^0+\epsilon$ and $s^0-\epsilon$. The players are unable to exploit either project in period 1 and, thus, must explore two new projects. Because the distribution of s_p is unbounded, for any δ, there exists a realization of s_p large enough such that $h(s^0+\epsilon, s_p) > \tilde{s}$. Finally, in this region (i.e., $h(s^0+\epsilon, s_p) > \tilde{s}$), the players follow the project selection rule of the symmetric-benefits benchmark and thus permanently exploit both projects. Therefore, with positive probability, the players exploit a project they have previously chosen not to exploit.

Statement 1 Statement 2 implies Statement 1.

Statement 2 Suppose $\delta \geq 1/3$ and $c = 1$, which ensures that $\mathcal{C}(\hat{s}_1, \hat{s}_2) \geq 1 = c\,\forall \hat{s}_1, \hat{s}_2$. When $\mathcal{C}(\hat{s}_1, \hat{s}_2) \geq 1$ but $h(\hat{s}_1, \hat{s}_2) < \tilde{s}$, the value of the second best project is irrelevant because within this range, the second best project will never be exploited since at most one project can be exploited. One can write the Bellman equation for the players: $B(\hat{s}_1, \hat{s}_2) = B(\max\{\hat{s}_1, \hat{s}_2\})$ when $h(\hat{s}_1, \hat{s}_2) < \tilde{s}$.

The indifference condition defining s^0 and the Bellman equation, $B^0(\cdot)$, corresponding to the symmetric-benefits benchmark is:

$$s^0 - 2 + \delta B^0(s^0) = 3 - 2 + \delta \mathbb{E}(B^0(\max\{s, s^0\})). \tag{13}$$

The left-hand side corresponds to the players' surplus when exploiting a project with value s^0. The right-hand side corresponds to the players' surplus when exploring one more project.

Suppose by contradiction, that the players never exploit a project with value less than s^0. In other words, suppose that the players weakly prefer to explore two new projects when the best project found so far is worth s^0:

$$\begin{aligned} & s^0 - 2 + (3-2) + \delta\mathbb{E}(B(\max\{s, s^0\}) + \delta\epsilon_1(\lambda) \\ \leq & 2(3-2) + \delta\mathbb{E}(B(\max\{s, s', s^0\}) + \delta\epsilon_2(\lambda). \end{aligned} \tag{14}$$

The first line corresponds to the value of exploiting a project worth s^0 and exploring an additional project. Under such a project selection rule, the first two terms correspond to the players' expected surplus in the current period and the latter two terms correspond to the continuation value. The term ϵ_1 corresponds to the change in continuation value upon finding a project valuable enough that $h(\hat{s}_1, \hat{s}_2) \geq \tilde{s}$. Specifically, $\epsilon_1(\lambda)$ corresponds to the probability that the new project's value, s', is sufficiently large such that $h(\hat{s}_1, \hat{s}_2) \geq \tilde{s}$, multiplied by the difference in continuation value in this region, as opposed to the continuation value when the continuation value is less than 2. The second line corresponds to the players' surplus following the exploration of two projects, where $\epsilon_2(\lambda)$ is defined analogously. Both ϵ_1, ϵ_2 approach 0 uniformly as $\lambda \to \infty$. These convergences happen because, to reach $h(\hat{s}_1, \hat{s}_2) \geq \tilde{s}$, the players must draw a project with value equal to at least 4. Because (i) drawing such a project occurs with probability approaching 0 as $\lambda \to \infty$ and (ii) the surplus differences associated with $\epsilon_1(\cdot)$ and$\epsilon_2(\cdot)$ remain bounded as $\lambda \to \infty$, the ϵ_1 and ϵ_2 terms uniformly decrease. One can then subtract Equation (13) from Inequality (14) and simplify using the closed-form solution of B^0 to derive:

$$\frac{1}{\lambda(1-\delta)} \leq \mathbb{E}(B(\max\{s, s', s^0\} - B(\max\{s, s^0\})) + \epsilon_2(\lambda) - \epsilon_1(\lambda). \tag{15}$$

Next, we show that $B(x) - B(y) \leq (x-y) \,/\, (1-\delta) + \epsilon_3(\lambda)$ when $x > y$, where ϵ_3 is exponentially decreasing in λ. Except for the exponentially decreasing probability that the players' continuation value exceeds two (which is accounted for by $\epsilon_3(\lambda)$), the players will be able to exploit at most one project per period. As such, the largest possible benefit from exploiting a better project occurs from exploiting the better project in every period. Utilizing such a bound, we can derive the following inequality:

$$\frac{1}{\lambda(1-\delta)} \leq \mathbb{E}\left(\frac{\max\{s, s', s^0\} - \max\{s, s^0\}}{1-\delta}\right) + \epsilon_2(\lambda) - \epsilon_1(\lambda) + \epsilon_3(\lambda) \tag{16}$$

$$\iff \frac{1}{\lambda(1-\delta)} \leq \frac{1}{1-\delta}\left(\frac{2}{\lambda} - \frac{1}{2\lambda} - \frac{1}{\lambda}\right) + \epsilon_2(\lambda) - \epsilon_1(\lambda) + \epsilon_3(\lambda). \tag{17}$$

Finally, because $\epsilon_1(\lambda), \epsilon_2(\lambda), \epsilon_3(\lambda)$ are all exponentially decreasing in λ, we can ignore these terms in the limit. Thus, one can further simplify to derive:

$$1 \leq \left(2 - \frac{1}{2} - 1\right), \tag{18}$$

which is a contradiction. $\square$