Multi-Project Collaborations
Online Appendix

Charles Angelucci and Roi Orzach*

February 9, 2025

1 Complete Characterization with Binary Support

Let $m = 2$ and consider a binary support distribution: $v_p \in \{\underline{v}, \bar{v}\}$ where $\bar{v} > \max\{\underline{v}, \tilde{v}\}$. We say a project is "suitable for exploitation" if and only if it has value $\bar{v}$. One can check that if $\alpha = \frac{1}{2}$ the players can always implement the first-best. For conciseness, we therefore assume $\alpha = 1$.

Proposition A1 *When $m = 2$ there are three possible non-empty experimentation policies which are optimal for a non-empty region of parameter values:*

1. **First Best:** *The players' scope of experimentation is maximal in all periods and the players exploit projects suitable for exploitation.*

2. **Immediate Exploitation:** *The players explore projects in one domain and, upon finding a project suitable for exploitation, they permanently exploit said project and immediately begin the exploration of projects in the other domain until identifying one suitable for exploitation.*

3. **Delayed Exploitation:** *The players explore projects in one domain and, upon finding a project suitable for exploitation, they idle on that domain and explore projects on the other domain until finding a project suitable for exploitation. Subsequently the players permanently exploit the two projects they have found which are suitable for exploitation.*

---
*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).

The players always prefer (1) to (2) and (2) to (3). For this reason, this proposition fully characterizes the optimal experimentation policy, since given any set of parameter values one can check which of these three policies is feasible.

Proof. The proof proceeds in two steps. First, we argue that no other experimentation policy may be optimal. Next, we use backwards induction to verify that these three policies are indeed optimal for a non-empty region of parameter values.

By symmetry, we can consider the experimentation strategy as a function of the number of projects suitable for exploitation, n. We first prove two results: scope is increasing with respect to n and scope is weakly larger than n.

To prove the former, note that any project selection rule that the players can implement when they have identified $n-1$ projects suitable for exploitation is also implementable when they have identified n such projects. This occurs because the continuation value of their relationship is strictly increasing in the number of projects suitable for exploitation. Additionally, exploiting a project suitable for exploitation results in a positive net present value. Hence, the scope of the players' relationship can never shrink on path.

The proof for the second statement notes that (i) by assumption, the players can always exploit n projects suitable for exploitation, (ii) the continuation value following exploration is weakly higher than that of exploitation, and (iii) the players maximize the scope as $\mathbb{E}(v_p) \geq 2c$. Combining these observations implies that the players' scope of experimentation always exceeds n.

Next, note that the players will neither exploit a project which is not suitable for exploitation (exploration provides higher current period profits and relaxes the implementability constraint), nor explore on a domain for which they have discovered a project suitable for exploitation (exploitation delivers higher current period profits, and both exploration and exploitation lead to an identical implementability constraint). Combining these observations with the result that scope cannot decrease implies that, if the players begin maximally, they can implement the first best. Therefore, it suffices to analyze experimentation policies where $|\mathbf{P}^1| = 1$, and to determine what the experimentation strategy dictates when $n = 1$. If when $n = 1$ the players maintain a scope of 1 and exploit this project, then the players never utilize the second domain. This cannot be optimal: by the logic of Bernheim and Whinston (1990), the players could instead implement the first best. If the players expand their scope, then this is necessarily the second policy in the proposition. If the players maintain

a scope of 1, then this is the third policy in the proposition.

For the second step of the proposition, note that the first-best is optimal as $\delta \to 1$. Finally, one can show that both (2) and (3) are optimal by computation: for instance, upon setting $\underline{v} = 0, \bar{v} = 4, c = 1, \frac{\delta}{1-\delta}(\bar{v} - 2c) = 1.01c$, there will exist two different values of $P(v_p = \bar{v})$ for which (2) and (3) are optimal for an open set of parameter values around these points.[1]

$\square$

2   Extensions

In this section we provide the proofs for the extensions outlined in Section 5.3 of the text.

2.1   Risky Collaborations

We make the following simplifications: $\alpha = 1$, $\mathbb{E}(v_p) = 2c$, and $m = 2$. Let the distribution of benefits on domain 1 be as in the main analysis. On domain 2 there exists only one project with value $v_r \in \{0, v\}$, where $\Pr(v_r = v) = q$ and $v \geq 2\tilde{v}$, ensuring that upon discovery of a project with value v the players can implement the first-best policy on domain 1.

Proposition A2 *For an open set of parameter values: (i) the players explore the domain 2 project after having first discovered a project with value exceeding $\tilde{v}$ on domain 1, (ii) despite being able to explore both domains in period 1.*

Proof of Proposition A2. We now provide the necessary conditions for the statement in the proposition to hold:

$$c \leq q\frac{\delta}{1 - \delta}(v - 2c) \tag{1}$$

$$c > \frac{\Pr(v_p \geq \tilde{v})\delta}{1 - \delta \cdot \Pr(v_p < \tilde{v})} \frac{1}{1 - \delta}\mathbb{E}(v_p - 2c|v_p \geq \tilde{v}) \tag{2}$$

$$c < \frac{\Pr(v_p \geq \tilde{v})\delta}{1 - \delta \cdot \Pr(v_p < \tilde{v})} \left(\frac{1}{1 - \delta}\mathbb{E}(v_p - 2c|v_p \geq \tilde{v}) + qv - 2c + q\frac{\delta}{1 - \delta}(v - 2c) \right) \tag{3}$$

---

[1]One could also instead prove that (2) is always implementable for a strictly larger region of discount factors than (1) to show that (2) is optimal.

These conditions state, respectively, that domain 2 could be started on its own, that domain 1 cannot be started on its own (or, equivalently, with knowledge that the domain 2 project is worth zero), and that domain 1 can be started when anticipating that domain 2 is started after discovering a project with value at least $\tilde{v}$. These conditions imply that the optimal experimentation policy takes one of two forms: (1) Explore projects on both domains, and if the risky project fails, then permanently idle on domain 1 if a project with value $\tilde{v}$ is not discovered in the first period. If the risky project has value v, then the experimentation policy becomes efficient on domain 1. (2) Explore domain 1 and only explore domain 2 once finding a project with value at least $\tilde{v}$. As policies (1) and (2) are both feasible, the players choose the policy with a higher net present value. Exploring both domains implies a surplus of:

$$
qv - 2c + q\frac{\delta}{1-\delta}(v - 2c)
$$
$$
+ (1-q)\frac{\delta}{1-\delta}\mathbb{E}((v_p - 2c)\mathbb{1}_{v_p \geq \tilde{v}}) + q\frac{\Pr(v_p > v^0)\delta}{1-\delta \cdot \Pr(v_p < \tilde{v})}\frac{1}{1-\delta}\mathbb{E}(v_p - 2c|v_p \geq v^0), \quad (4)
$$

where the first line corresponds to the surplus from domain 2 and the second line corresponds to the surplus from domain 1. Further, the surplus from exploring domain 1 and later exploring domain 2 is exactly Equation (3) as the expected surplus in period 1 is equal to zero. Let us hold fixed the surplus from domain 2 and δ, but let us consider varying q while v remains fixed to maintain the same expected surplus. When $q = 1$, there is no uncertainty and the surplus from exploring both domains at once can be shown to be strictly greater than delaying. When domain 2 is explored after domain 1, changing q but holding the expected surplus from domain 2 fixed has no effect on the total expected surplus. In contrast, when exploring both domains at once, a decrease in q increases the probability that domain 1 is never explored. As q continues to decrease, the profit from exploring both domains at once converges to the expected surplus from domain 2 plus the probability that the first drawn project in domain 1 exceeds $\tilde{v}$, multiplied by the profitability of exploiting such a project. In contrast, the profit from going sequentially converges to that of with certainty utilizing domain 1 plus the expected profitability of domain 2 when explored after finding a suitable project in domain 1. One can see that the players may delay domain 2, even if it is more profitable than domain 1, so as to prevent permanently idling on domain 1. □

4

2.2 Symmetric Benefits

Let us first consider a case where on one domain the benefits are symmetric but on the other domain benefits are asymmetric. Namely, for domain 1 all projects equally benefit both parties but domain 2 is as in our main analysis.

Proposition A3 *Let δ^* and $\bar{\delta}$ be defined as in the main analysis. In this extension, $\delta^* < \bar{\delta}$. In other words, for intermediate discount factors, the scope of experimentation is initially limited on path, with scope increasing with strictly positive probability along the equilibrium path.*

Proof of Proposition A3. Observe that $\delta^* = 0$ because an experimentation policy consisting of exploration on the symmetric benefits domain and only expanding the scope once a project exceeding $\tilde{v}$ is discovered is always feasible. Further, $\bar{\delta} > 0$ as $\mathcal{C}(\cdot) \to 0$ as $\delta \to 0$, implying that the players cannot explore the asymmetric benefits domain in period 1. $\qquad\square$

Let us now consider the case where α, the probability that the benefits go to player 1, differs across the domains. Denote by α_j, the corresponding probability for domain j. For simplicity, let us assume $\alpha_1 = 1/2$ and $\alpha_2 = 1$.

Proposition A4 *Let δ^* and $\bar{\delta}$ be defined as in the main analysis. In this extension, $\delta^* < \bar{\delta}$. In other words, for intermediate discount factors, the scope of experimentation is initially limited on path, with scope increasing with strictly positive probability along the equilibrium path.*

Proof of Proposition A4. Observe that $\delta^* = 0$ because an experimentation policy consisting of exploration on domain 1 and only expanding scope once a project exceeding $2\tilde{v}$ is discovered is always feasible. Further, $\bar{\delta} > 0$ because $\mathcal{C}(\cdot) \to 0$ as $\delta \to 0$, implying that the players cannot explore projects in domain 2 in period 1. $\qquad\square$

2.3 Technological Interdependencies

In this extension, we assume that domains are perfectly technologically dependent. Formally, assume there is a bijection between the projects in the two domains, such that for any two projects paired by the bijection, the project values are equal. Such an assumption implies that any project value found on a given domain can be used

to produce an identical value on the second domain. Further, for simplicity, we will again restrict attention to $\alpha = 1$ and $m = 2$, but the forces naturally extend.

Proposition A5 *Let δ^* and $\bar{\delta}$ be defined as in the main analysis. In this extension, $\delta^* < \bar{\delta}$. Further, scope is terminally maximal with probability 1 for any non-empty experimentation.*

Proof of Proposition A5. The latter statement follows directly from technological interdependencies and Bernheim and Whinston (1990): if the players permanently exploit k projects with value $\hat{v}$, then the players can exploit m projects with value $\hat{v}$.

In this analysis, Proposition 1 continues to hold and the optimal experimentation policy only conditions on the best project found across all domains: $\hat{v}$. Denote by $B(\hat{v})$ the bellman equation corresponding to the players' strategy. Therefore, $\bar{\delta}$ is defined as

$$2c = \bar{\delta}\mathbb{E}_{\hat{v}_1,\hat{v}_2}(B(\max\{\hat{v}_1,\hat{v}_2\})|\bar{\delta}). \tag{5}$$

To prove, $\delta^* < \bar{\delta}$, it suffices to show:

$$c < \bar{\delta}\lim_{\delta \uparrow \bar{\delta}} \mathbb{E}_{\hat{v}_1}(B(\hat{v}_1))|\delta) \iff 2\lim_{\delta \uparrow \bar{\delta}} \mathbb{E}_{\hat{v}_1}(B(\hat{v})|\delta) > \mathbb{E}_{\hat{v}_1,\hat{v}_2}(B(\max\{\hat{v}_1,\hat{v}_2\})|\bar{\delta}). \tag{6}$$

$B(\cdot|\bar{\delta})$ is characterized by two thresholds, $\underline{\hat{v}}, \bar{\hat{v}}$: for $\hat{v} \geq \bar{\hat{v}}$ the players exploit two projects with value $\hat{v}$, for $\hat{v} \in [\underline{\hat{v}}, \bar{\hat{v}})$ the players exploit one project with value $\hat{v}$ and explore a new project, and for $\hat{v} < \underline{\hat{v}}$ the players explore two new projects.

Consider an alternative experimentation strategy with $\epsilon > 0$: for $\hat{v} \geq \bar{\hat{v}} + \epsilon$ the players exploit two projects with value $\hat{v}$, for $\hat{v} \in [\underline{\hat{v}} + \epsilon, \bar{\hat{v}} + \epsilon)$ the players exploit one project with value v and explore a new project, and for $\hat{v} < \underline{\hat{v}} + \epsilon$ the players explore one new project. As $\epsilon > 0$, this experimentation policy is implementable for $\hat{v} > \underline{\hat{v}} + \epsilon$ as $\delta \uparrow \bar{\delta}$. Hence, it suffices to show that this policy is implementable at period 1.

Under this policy, $\lim_{\delta \uparrow \bar{\delta}} B(0|\delta) > \frac{1}{2}B(0|\bar{\delta})$. Therefore, the expected continuation value of this policy at date 1 when $\delta \uparrow \bar{\delta}$ is strictly greater than half the expected continuation value of the optimal policy when $\delta = \bar{\delta}$. This assertion is exactly the claim we needed to show Equation (6) holds.

$\square$

References

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