

CeNDEF Working Paper 26–06

Grega Smrkolj and Florian Wagener

R&D Competition and Cooperation with Distance-Dependent Spillovers

October 2026

Abstract

We study a continuous-time duopoly model of process innovation with R&D spillovers, comparing noncooperative R&D with cooperative research regimes. We extend the standard constant-spillover framework by allowing knowledge transmission to decay with technological distance and to favor followers over leaders in asymmetric specifications. In a global Markov-perfect model, firms may invest before production is viable, enter or exit production as costs evolve, and converge to no-market, monopoly, or duopoly outcomes. State-dependent spillovers change R&D incentives, catch-up dynamics, long-run market structure, and the welfare effects of research cooperation. In the computed equilibria, more follower-favoring spillovers weaken the leader's private incentive to invest but accelerate catch-up, shorten monopoly phases, and make eventual duopoly more likely. When spillovers are weak, cooperation mainly softens dynamic rivalry; when information sharing is substantial, cooperation expands market formation, lowers long-run costs, and can raise both consumer and total surplus, especially under the research-joint-venture regime. The value of R&D cooperation depends on the direction and persistence of knowledge flows, not only on their average intensity.



R&D Competition and Cooperation with Distance-Dependent Spillovers

Grega Smrkolj*

Florian Wagener†

June 2026

Abstract

We study a continuous-time duopoly model of process innovation with R&D spillovers, comparing noncooperative R&D with cooperative research regimes. We extend the standard constant-spillover framework by allowing knowledge transmission to decay with technological distance and to favor followers over leaders in asymmetric specifications. In a global Markov-perfect model, firms may invest before production is viable, enter or exit production as costs evolve, and converge to no-market, monopoly, or duopoly outcomes. State-dependent spillovers change R&D incentives, catch-up dynamics, long-run market structure, and the welfare effects of research cooperation. In the computed equilibria, more follower-favoring spillovers weaken the leader's private incentive to invest but accelerate catch-up, shorten monopoly phases, and make eventual duopoly more likely. When spillovers are weak, cooperation mainly softens dynamic rivalry; when information sharing is substantial, cooperation expands market formation, lowers long-run costs, and can raise both consumer and total surplus, especially under the research-joint-venture regime. The value of R&D cooperation depends on the direction and persistence of knowledge flows, not only on their average intensity.

Keywords: R&D cooperation, spillovers, process innovation, dynamic oligopoly, welfare

JEL: C63, C73, D43, D92, L13, O31, O32

1 Introduction

Hardly any commercially valuable idea is immune to imitation. Knowledge can cross firm boundaries through channels such as patent disclosure, employee mobility, reverse engineering, and informal contacts, and may also be transmitted deliberately through licensing and research collaboration. Because these flows are not always fully priced or controlled, firms may fail to appropriate the full value of their own R&D. In models of process innovation, such spillovers

*Corresponding author. Newcastle University Business School, Newcastle University, 5 Barrack Road, Newcastle upon Tyne NE1 4SE, UK. E-mail: grega.smrkolj@newcastle.ac.uk.

†CeNDEF, Amsterdam School of Economics, University of Amsterdam, Roetersstraat 11, 1018 WB Amsterdam, The Netherlands; Tinbergen Institute; CESifo. E-mail: f.o.wagener@uva.nl.

are commonly represented by allowing one firm's R&D to reduce, at least in part, the rival's production cost. Empirically, these flows are large: using U.S. firm-level data, [Bloom et al. \(2013\)](#) estimate that the social returns to R&D are at least twice as high as the private returns. The gap between private and social returns is the classic reason why privately chosen R&D may differ from the socially desirable level.

The same logic underlies the relatively permissive treatment of R&D cooperation in competition policy.¹ These policies reflect the same basic idea: cooperation in research can help internalize knowledge spillovers. Yet in dynamic industries, firms' relative technological positions evolve endogenously, and the strength and direction of knowledge flows may change with them. This raises the central question of the paper: when does R&D cooperation improve welfare, and when does it mainly soften dynamic rivalry? We answer this question in a dynamic model in which spillovers decay with technological distance and may favor followers over leaders.

The interplay between appropriability, knowledge diffusion, and competitive incentives has made R&D spillovers and research cooperation central topics in industrial organization. The seminal theoretical contribution is [d'Aspremont and Jacquemin \(1988\)](#), who study a two-stage model of process innovation with spillovers. Firms first choose cost-reducing R&D and then compete in the product market. When spillovers are low, firms appropriate most of the benefits of their own R&D, so cooperation has little knowledge externality to internalize. When spillovers are high, noncooperative incentives are weakened by free riding, and cooperation becomes valuable because it internalizes large knowledge flows. In this setting, R&D cooperation improves welfare when spillovers are sufficiently large. [Kamien et al. \(1992\)](#) show that the effects of cooperation also depend on the distinction between coordinating R&D choices and sharing research outcomes. In their framework, the research-joint-venture (RJV) regime performs best because it combines joint R&D decisions with full information sharing.

These early models generated numerous extensions and generalizations. [De Bondt \(1997\)](#) surveys this literature and concludes that spillovers have ambiguous effects: they weaken appropriability but transmit useful knowledge, so innovative output is often highest at intermediate levels of appropriation. Subsequent work shows that these conclusions depend on market structure, firm asymmetries, uncertainty, equilibrium selection, and the way spillovers and appropriability enter the innovation process. [Suzumura \(1992\)](#) shows that the direction of R&D distortions depends on spillovers: high spillovers can leave both noncooperative and cooperative R&D below even second-best levels, whereas without spillovers noncooperative R&D may exceed social benchmarks. [Amir and Wooders \(1998\)](#) show that asymmetric noncooperative R&D equilibria can arise and may overturn the usual ranking of competition and cooperation. [Salant and Shaffer \(1998\)](#) show that identical firms in a research joint venture may choose asymmetric R&D investments to raise joint profits, and that this asymmetry can raise welfare even without

¹For an overview of the competition-policy treatment of cooperative R&D, see [Motta \(2004\)](#). In the United States, the National Cooperative Research Act of 1984, later amended as the National Cooperative Research and Production Act, limited antitrust exposure for qualifying research joint ventures. In the European Union, qualifying R&D agreements may fall under the research and development block exemption regulation. The EU also directly promotes and funds collaborative research through its research and innovation framework programmes, including Horizon Europe.

spillovers. [Martin \(2002\)](#) distinguishes input spillovers from post-innovation appropriability in a stochastic innovation model and shows that welfare is favored by high input spillovers and low appropriability, with joint R&D improving welfare by increasing effective spillovers and sustaining post-innovation rivalry.

Another strand endogenizes spillovers. Rather than treating knowledge transmission as fixed, these papers make absorption or disclosure a strategic choice. They show that cooperation does not automatically imply maximal information sharing, so its welfare case depends on how it changes knowledge transmission (see, among others, [Katsoulacos and Ulph, 1998](#); [Kamien and Zang, 2000](#); [Poyago-Theotoky, 1999](#); [Leahy and Neary, 2007](#); [Amir et al., 2003](#)). This literature is close in spirit to the present paper because it treats knowledge transmission as endogenous rather than fixed, but it is largely static.

Dynamic models of R&D with spillovers are fewer. In a discrete-time framework, [Petit and Tolwinski \(1999\)](#) study R&D cooperation and competition, including asymmetric firms, and show that research joint ventures are more likely to improve welfare when spillovers are substantial, whereas competition tends to perform better when spillovers are limited. In a stochastic discrete-time model of process innovation, [Samano et al. \(2017\)](#) similarly find that the welfare case for R&D cooperation strengthens with spillovers: long-run welfare is highest under the RJV regime, and consumer-surplus gains from cooperation increase with spillovers.

[Cellini and Lambertini \(2009\)](#) introduce a continuous-time version of the d'Aspremont–Jacquemin framework in which R&D is an investment flow that drives nonlinear dynamics in firms' marginal costs. Starting from given initial cost levels, firms choose R&D continuously as costs converge gradually to their steady-state values. Building on this continuous-time approach, [Hinlopen et al. \(2013\)](#) develop a global dynamic framework of process innovation in which initial conditions can determine qualitatively different long-run outcomes. In particular, the initial marginal cost determines not only how much a firm invests, but also whether a technology is developed at all. R&D and production need not coincide: a firm may invest before production becomes viable, produce while continuing to innovate, or never bring a technology to market. Related dynamic work has also increasingly emphasized the role of initial conditions in shaping long-run outcomes. For example, in [Dawid et al. \(2015\)](#), an incumbent's incentive to invest in product innovation depends on its initial product capacity and knowledge stock. [Smrkolj and Wagener \(2019\)](#) extend the global framework to a stochastic dynamic game with two competing firms. The analysis shows how spillovers affect R&D in existing markets, the emergence of new markets, and long-run market structure: production may ultimately be carried out by one firm, by both firms, or by neither.

The present paper builds on this global dynamic framework and shares with [Samano et al. \(2017\)](#) the aim of comparing competitive and cooperative R&D regimes in process innovation. It differs in two respects. First, the comparison is embedded in a continuous-time global feedback model with different cost dynamics, in which market formation and long-run market structure arise endogenously as firms' costs evolve. Firms may begin R&D before production is viable, enter or leave production as costs cross market-entry boundaries, and end up in no-market, monopoly, or duopoly outcomes. Second, spillovers are state dependent. Whereas [Samano et al. \(2017\)](#) study

fixed spillover probabilities, knowledge transmission here decays with technological distance and may favor followers over leaders, reflecting differences in the strength and value of knowledge flows across firms' relative technological positions. This creates a feedback loop: R&D changes the cost gap, the cost gap changes spillover intensity, and spillover intensity changes the marginal returns to own and rival R&D. As a result, the level, direction, and persistence of spillovers evolve along the industry's transition path and are central to catch-up, market formation, long-run market structure, and the welfare effects of R&D cooperation. Because cooperation coordinates R&D incentives without eliminating firm-level state variables, it can also affect how innovative activity and production are allocated across firms. Even in the research-joint-venture regime, where full information sharing equalizes the deterministic knowledge-flow component, firm-specific uncertainty continues to generate idiosyncratic cost movements.

The empirical evidence also points to spillovers that vary with firms' technological positions. Already in [Jaffe \(1986\)](#), outside R&D raises the productivity of firms' own R&D when it is conducted by technological neighbors. [Orlando \(2004\)](#) sharpen this point by showing that industrial R&D spillovers are strongest among technologically similar firms. Spillovers are attenuated by technological distance, whereas geographic distance does not significantly attenuate spillovers within narrowly defined technological groups. [Bloom et al. \(2013\)](#) show that R&D by technologically close firms is associated with higher market value, patenting, and productivity. Taken together, these findings suggest that technological closeness facilitates the absorption and use of external R&D. They motivate modelling spillover intensity as declining with technological distance.

Technological distance, however, need not affect leaders and followers symmetrically. External knowledge is not equally valuable to all recipients. A lagging firm can often use knowledge generated closer to the frontier to close a technological gap; a leader has less to learn from a follower whose knowledge is less advanced or already familiar. This logic is consistent with evidence that imitation is most productive away from the frontier. [Liao \(2020\)](#) finds that imitation raises productivity more for laggard firms than for leaders, while the productivity effect of innovation increases as firms approach the frontier. The same logic appears in the distance-to-frontier literature: imitation is relatively more important away from the technology frontier, whereas innovation becomes relatively more important closer to the frontier ([Acemoglu et al., 2006](#)). Firm size may proxy for technological position in some environments: [Acs et al. \(1994\)](#) find that R&D generated by large firms and universities contributes more to the innovative activity of small firms than to that of large firms. Evidence on location choices points in the same direction. [Alcácer and Chung \(2007\)](#) argue that advanced firms are potential spillover sources for less advanced competitors and find that less technologically advanced firms favor locations with high industrial innovative activity, while technologically advanced firms avoid such locations, consistent with laggards seeking knowledge inflows and leaders limiting knowledge outflows. [Livannis and Lamin \(2016\)](#) similarly find that technological laggards' R&D efforts are aided by the presence of nearby leaders, partly because laggards look to such leaders as sources of skilled workers. [Griffith et al. \(2004\)](#) show that R&D also builds absorptive capacity, and that this channel is quantitatively more important farther from the technological frontier. Together, these findings motivate a specification in which spillovers decline with technological distance but do

so more slowly for followers than for leaders.

Our analysis delivers three main insights. First, the welfare effect of R&D cooperation is state dependent. When spillovers are weak, cooperation has little knowledge externality to internalize and mainly softens dynamic rivalry. R&D falls over much of the duopoly region, cost asymmetries widen, and the industry is more likely to move toward production by a single firm. With distance-decaying spillovers, cooperation is most valuable when firms are technologically close, because knowledge flows are then strongest and coordinated R&D generates common cost reductions. As firms move apart, these knowledge flows weaken and the efficiency case for cooperation becomes less robust. The same tradeoff appears across cost levels: cooperation is most attractive near the development margin, where coordinated R&D can make production viable, but may reduce welfare when costs are already low and its main effect is to soften rivalry rather than generate substantial additional cost savings.

Second, the direction of spillovers matters. More follower-favoring spillovers weaken the leader's private incentive to invest, but they also direct knowledge flows toward the firm for which they have the highest marginal value. In the computed equilibria, this second force dominates. Stronger follower-favoring spillovers accelerate catch-up, shorten monopoly phases, enlarge the set of initial conditions from which both firms remain active, and make eventual duopoly more likely. Thus, spillover asymmetry matters not only for R&D incentives at a given state, but also for endogenous transition dynamics and long-run market structure.

Third, the value of cooperation depends on information sharing. R&D cartelization alone can reduce investment and preserve asymmetries when spillovers are weak. By contrast, the research-joint-venture regime combines coordinated R&D decisions with full information sharing. Among the cooperative regimes considered, it provides the strongest dynamic-efficiency case for cooperation: it reduces the no-market region, lowers long-run costs, and generates the most robust gains in consumer surplus and total surplus. Even this regime, however, does not dominate noncooperative R&D in every respect. In parts of the state space, especially at lower cost levels and with asymmetric firms, noncooperative R&D can preserve duopoly more effectively.

These results matter for policy because R&D cooperation is often justified by the need to internalize knowledge spillovers. We show that this rationale depends on the evolution of knowledge flows, not only on their average intensity. A fixed average spillover can overstate or understate the gains from cooperation when spillovers vary with technological distance, and the direction of knowledge flows can change transition dynamics and long-run market structure. In particular, R&D cooperation can induce monopolization even without explicit product-market collusion. Policies that encourage or evaluate collaborative research should therefore account for both the state dependence of spillovers and their direction.

The nonlinear Markov-perfect equilibrium problem generated by our model does not yield a tractable closed-form solution, so we solve it numerically using a monotone semi-Lagrangian scheme for the Hamilton–Jacobi–Bellman equations. This method allows us to compute global feedback policies over the relevant state space and to compare R&D regimes from arbitrary initial cost configurations, including asymmetric ones.

The rest of the paper is organized as follows. Section 2 presents the model, including the spillover specifications and R&D regimes. Section 3 describes the numerical solution method. Section 4 reports the numerical results, first for alternative spillover specifications under noncooperative R&D and then for cooperative R&D regimes and welfare. Section 5 concludes. The Appendix reports robustness checks.

2 The model

Time is continuous, $t \in [0, \infty)$. Two firms, $i \in \{1, 2\}$, compete in a Cournot product market and have linear production costs $C_i(q_i(t)) = c_i(t) q_i(t)$, where $c_i(t)$ is firm i 's marginal cost. At each t , firms simultaneously choose output $q_i(t)$ and R&D effort $k_i(t)$. R&D reduces marginal costs over time. We use $-i$ to denote firm i 's rival and suppress time arguments whenever no confusion arises.

2.1 R&D technology and spillovers

Firm i chooses R&D effort $k_i(t) \geq 0$ at instantaneous cost $bk_i(t)^2$, $b > 0$. Effective R&D is given by

$$K_i(t) = k_i(t) + \beta(\cdot) k_{-i}(t), \quad (1)$$

where $k_{-i}(t)$ denotes the rival's R&D effort and the spillover function $\beta(\cdot) \in [0, 1]$ determines the share of the rival's knowledge that benefits firm i . The case $\beta(\cdot) = 0$ corresponds to no spillovers, while $\beta(\cdot) = 1$ corresponds to full spillovers.

We measure technological distance in relative terms, so that the technology gap is invariant to the overall cost level and to the units in which costs are measured. Let $x_i = \log c_i$ denote firm i 's log marginal cost. We define the (signed) technology gap as

$$\Delta x_i \equiv x_i - x_{-i} = \log \left(\frac{c_i}{c_{-i}} \right). \quad (2)$$

Firm i is a *follower* if $\Delta x_i > 0$ (equivalently $c_i > c_{-i}$) and a *leader* otherwise.

We consider three alternative specifications of $\beta(\cdot)$, which differ in whether spillovers are exogenous or depend on technological distance:

(i) **Exogenous spillovers (baseline).**

$$\beta(\cdot) = \bar{\beta}, \quad \bar{\beta} \in [0, 1]. \quad (3)$$

Under this specification, each firm benefits from a fixed fraction $\bar{\beta}$ of the rival's R&D effort, irrespective of technological distance. This is the standard constant-spillover benchmark.

(ii) **Symmetric distance-decaying spillovers.**

To capture the empirically documented idea that firms benefit most from technologically

similar rivals, we allow spillovers to decline with the log-cost gap Δx_i . We set

$$\beta(\Delta x_i) = \beta_L + (\beta_H - \beta_L)e^{-\kappa|\Delta x_i|}, \quad (4)$$

where $0 \leq \beta_L < \beta_H \leq 1$ and $\kappa > 0$. Spillovers peak at parity $\Delta x_i = 0$ with value β_H and decline exponentially toward β_L as $|\Delta x_i|$ increases. The parameter κ governs the rate of decay (see Fig. 1a).²

(iii) **Asymmetric distance-decaying spillovers.**

Empirical evidence further suggests that followers benefit more from leaders than vice versa. We capture this by allowing spillovers to decay with technological distance at different rates depending on whether firm i is the leader or the follower. We set

$$\beta(\Delta x_i) = \begin{cases} \beta_L + (\beta_H - \beta_L)e^{-\kappa|\Delta x_i|}, & \Delta x_i \leq 0, \\ \beta_L + (\beta_H - \beta_L)e^{-\alpha\kappa|\Delta x_i|}, & \Delta x_i > 0, \end{cases} \quad \alpha \in (0, 1], \kappa > 0. \quad (5)$$

Spillovers again peak at parity, $\Delta x_i = 0$, with value β_H and decline toward β_L as the cost gap widens. The parameter κ governs the baseline rate of decay, while α scales the follower-side decay rate, $\kappa_F = \alpha\kappa$ (see Fig. 1b). Thus, for $\alpha < 1$, spillovers decay more slowly for the follower, implying $\beta(\Delta x) > \beta(-\Delta x)$ for all $\Delta x > 0$. We therefore refer to lower values of α as more follower-favoring spillovers. Setting $\alpha = 1$ nests the symmetric distance-decaying specification in (4).³

Marginal costs (c_1, c_2) are the state variables. Their dynamics is described by a stochastic cost process

$$\frac{dc_i}{c_i} = (\delta - K_i) dt + \sigma dB_i, \quad (6)$$

a geometric Itô diffusion with initial condition $c_i(0) = c_{i,0} > 0$. The constant drift parameter $\delta > 0$ reflects efficiency reduction due to the aging of technology and organizational forgetting, so that maintaining efficiency requires ongoing investment. Effective R&D K_i reduces firm i 's marginal cost. The diffusion term σdB_i , where the B_i are independent standard Brownian motions and $\sigma > 0$ is a volatility parameter, captures random discoveries as well as operational and input-cost disturbances that affect production efficiency. This formulation has two economically plausible features. First, marginal costs remain strictly positive (if $c_{i,0} > 0$, then $c_i(t) > 0$ for all $t \geq 0$) and cost reductions are proportional rather than additive. Since the drift in (6) scales with the current cost level, a given effective R&D intensity K_i yields larger absolute cost reductions when c_i is high than when c_i is low. Thus, lowering costs from a high level is easier

²The parameter κ controls how rapidly spillovers fall from β_H toward the lower bound β_L . In particular, $\beta(\Delta x)$ equals the midpoint $(\beta_H + \beta_L)/2$ at $|\Delta x| = \ln 2/\kappa$, i.e. when the follower's marginal cost is $2^{1/\kappa}$ times the leader's. Thus, for $\kappa = 1$ the midpoint is reached at a 2:1 cost ratio, whereas for $\kappa = 3$ at approximately 1.26:1.

³The piecewise formulation in (5) is equivalent to the compact expression

$$\beta(\Delta x_i) = \beta_L + (\beta_H - \beta_L) \exp \left[-\kappa|\Delta x_i| + (1 - \alpha)\kappa \max\{\Delta x_i, 0\} \right].$$




Figure 1: Distance-decaying spillovers. Panel (a) illustrates the symmetric specification (4) for two values of the decay parameter κ ; larger κ implies a faster decline of spillovers with technological distance. Panel (b) illustrates the leader–follower asymmetric specification (5): holding $\kappa = 1$ fixed, lower values of α reduce the decay rate of spillovers on the follower side ($\Delta x_i > 0$). In both panels, $\beta_L = 0.2$ and $\beta_H = 0.8$.

than improving an already efficient technology. Second, random shocks scale with the cost level, so uncertainty is greater when the technology is less mature.

2.2 Product market

At each instant t , given current costs $(c_1(t), c_2(t))$, firms play a static Cournot game with inverse demand

$$p(t) = \max\{A - q_1(t) - q_2(t), 0\}, \quad (7)$$

where $A > 0$ is the demand intercept, or choke price. Because output choices do not affect the evolution of marginal costs, the product-market equilibrium at each time t is the static Cournot–Nash equilibrium conditional on the current state $(c_1(t), c_2(t))$.

If both firms are active ($q_1^* > 0$ and $q_2^* > 0$), equilibrium quantities are $q_i^*(c_i, c_{-i}) = (A - 2c_i + c_{-i})/3$, which requires $2c_i - c_{-i} < A$ for each i . If instead $2c_i - c_{-i} \geq A$, firm i is inactive ($q_i^* = 0$) and firm $-i$ supplies the monopoly quantity $q_{-i}^* = (A - c_{-i})/2$, provided that $c_{-i} < A$. If $c_1 \geq A$ and $c_2 \geq A$, no production occurs. These activity conditions define the static Cournot entry–exit curves in cost space, denoted by E_1 and E_2 , which separate states in which a firm is active in the product market from states in which it is inactive. They also imply the following piecewise sales-profit function, where the first line corresponds to interior duopoly, the second to monopoly by firm i , and the third to states in which firm i is inactive:

$$g(c_i, c_{-i}) = \begin{cases} \frac{(A - 2c_i + c_{-i})^2}{9}, & \text{if } 2c_i - c_{-i} < A \text{ and } 2c_{-i} - c_i < A, \\ \frac{(A - c_i)^2}{4}, & \text{if } c_i < A \text{ and } 2c_{-i} - c_i \geq A, \\ 0, & \text{otherwise.} \end{cases} \quad (8)$$

The (c_1, c_2) state space is thereby partitioned into no-production, duopoly, and monopoly regions, with the prevailing market structure determined by the firms' current cost positions. Even when production is not currently viable (for example, if $c_i \geq A$), a firm may still invest in R&D to reduce future costs and enter production once costs fall sufficiently, capturing technologies that require upfront investment before production becomes viable.

Instantaneous total profit equals sales profit minus R&D expenditure:

$$\pi_i = g(c_i, c_{-i}) - b k_i^2. \quad (9)$$

2.3 Optimal R&D investment and Markov strategies

We consider two R&D regimes. The benchmark regime is noncooperative R&D, in which each firm chooses its own R&D effort taking the rival's strategy as given. We then consider an R&D cartel, in which firms coordinate their R&D efforts so as to maximize joint discounted profits, while product-market competition remains unchanged.

2.3.1 Noncooperative R&D

We analyze *Markovian* or *feedback* R&D strategies, under which firm i 's R&D effort is a function of the current state, $k_i(t) = \Gamma_i(c(t)) = \Gamma_i(c_1(t), c_2(t))$, $i \in \{1, 2\}$. Feedback strategies are non-anticipative, that is, the R&D effort taken at time t is a function of the information available to the firm at that time.

Let $\rho > 0$ denote the discount rate. Given the rival's strategy Γ_2 , firm 1 chooses $k_1 = \Gamma_1(c)$ to maximize, for every initial state $c_0 = (c_{1,0}, c_{2,0})$, expected discounted profits

$$J^1(c_0, k_1, \Gamma_2) = \mathbb{E}_{c_0} \left[\int_0^\infty (g(c) - b k_1^2) e^{-\rho t} dt \right], \quad (10)$$

subject to the cost dynamics

$$dc = \begin{pmatrix} \left(\delta - k_1 - \beta \left(\log \frac{c_1}{c_2} \right) \Gamma_2(c) \right) c_1 \\ \left(\delta - \Gamma_2(c) - \beta \left(\log \frac{c_2}{c_1} \right) k_1 \right) c_2 \end{pmatrix} dt + \begin{pmatrix} \sigma c_1 & 0 \\ 0 & \sigma c_2 \end{pmatrix} dB. \quad (11)$$

The strategy Γ_1 is then a *best response* to the strategy Γ_2 . The problem of firm 2 is analogous.

A *Markov-perfect equilibrium* is a pair of feedback strategies (Γ_1^*, Γ_2^*) such that each Γ_i^* is a best response to Γ_{-i}^* .

Log-cost coordinates. For numerical convenience, we express the problem in log-cost coordinates $x = (x_1, x_2)$. We define the transformed sales profit as $h_i(x) \equiv g(e^{x_i}, e^{x_{-i}})$. Expected discounted profits of firm 1 are

$$J^1(x_0, k_1, \Gamma_2) = \mathbb{E}_{x_0} \left[\int_0^\infty (h_1(x) - b k_1^2) e^{-\rho t} dt \right], \quad (12)$$

given the dynamics (13) and the strategy choice Γ_2 of the other firm. By Itô's lemma, the log-cost dynamics are

$$dx = f(x, k_1, \Gamma_2(x)) dt + \sigma dB = \begin{pmatrix} \delta - \frac{1}{2}\sigma^2 - k_1 - \beta(x_1 - x_2)\Gamma_2(x) \\ \delta - \frac{1}{2}\sigma^2 - \Gamma_2(x) - \beta(x_2 - x_1)k_1 \end{pmatrix} dt + \sigma dB. \quad (13)$$

Again, the situation of firm 2 is analogous.

Hamilton–Jacobi–Bellman (HJB) system. Let (Γ_1^*, Γ_2^*) be a Markov-perfect equilibrium, and let $V^i(x) = V^i(x_1, x_2)$ denote firm i 's equilibrium value function in log-cost coordinates, that is

$$V^i(x_0) = J^i(x_0, \Gamma_1^*, \Gamma_2^*).$$

By the principle of optimality, we have that

$$V^i(x_0) = \max_{k_i \geq 0} \mathbb{E}_{x_0} \left[\int_0^{\Delta t} (h_i(x(t)) - bk_i^2) e^{-\rho t} dt + e^{-\rho \Delta t} V^i(x(\Delta t)) \right] \quad (14)$$

Let $\Delta V^i = \frac{\partial^2 V^i}{\partial x_1^2} + \frac{\partial^2 V^i}{\partial x_2^2}$ denote the Laplacian of V^i . In equilibrium, the value functions V^1 and V^2 solve the system of Hamilton–Jacobi–Bellman equations

$$\begin{aligned} \rho V^i(x) = \max_{k_i \geq 0} & \left[h_i(x) - bk_i^2 + \frac{\partial V^i}{\partial x_i}(x) \left(\delta - \frac{1}{2}\sigma^2 - k_i - \beta(x_i - x_{-i})\Gamma_{-i}^*(x) \right) \right. \\ & \left. + \frac{\partial V^i}{\partial x_{-i}}(x) \left(\delta - \frac{1}{2}\sigma^2 - \Gamma_{-i}^*(x) - \beta(x_{-i} - x_i)k_i \right) \right] + \frac{1}{2}\sigma^2 \Delta V^i(x), \end{aligned}$$

$i \in \{1, 2\}$. As this is a system of uniformly elliptic equations, its solutions are twice continuously differentiable functions. We finally note that the Hamilton–Jacobi–Bellman equation is obtained as the limit $\Delta t \rightarrow 0$ of the principle of optimality.

Define the marginal payoff to R&D as the derivative of the HJB maximand with respect to k_i ,

$$\Psi_i(k_i; x) \equiv -2bk_i - \frac{\partial V^i}{\partial x_i}(x) - \beta(x_{-i} - x_i) \frac{\partial V^i}{\partial x_{-i}}(x).$$

The optimal effort $k_i^* = \Gamma_i^*(x)$ satisfies the complementarity condition

$$k_i^* \geq 0, \quad \Psi_i(k_i^*; x) \leq 0, \quad k_i^*(x) \Psi_i(k_i^*; x) = 0. \quad (15)$$

Thus, if $k_i^* > 0$, then $\Psi_i(k_i^*; x) = 0$, which yields

$$\Gamma_i^*(x) = \max \left\{ 0, -\frac{\frac{\partial V^i}{\partial x_i}(x) + \beta(x_{-i} - x_i) \frac{\partial V^i}{\partial x_{-i}}(x)}{2b} \right\}.$$

2.3.2 R&D cooperation

We next consider an R&D cartel in which firms coordinate R&D efforts (k_1, k_2) to maximize joint discounted profits. Relative to the benchmark noncooperative regime, the cooperative regime differs only in the R&D decision rule. Firms continue to compete in the product market and to operate separate laboratories, so innovation outcomes and cost dynamics remain subject to firm-specific uncertainty.

Let $W(x)$ denote the cooperative value function. It satisfies

$$\begin{aligned} \rho W(x) = \max_{k_1, k_2 \geq 0} & \left[h(x_1, x_2) + h(x_2, x_1) - b(k_1^2 + k_2^2) \right. \\ & + \frac{\partial W}{\partial x_1}(x) \left(\delta - \frac{1}{2}\sigma^2 - k_1 - \beta(x_1 - x_2)k_2 \right) \\ & \left. + \frac{\partial W}{\partial x_2}(x) \left(\delta - \frac{1}{2}\sigma^2 - k_2 - \beta(x_2 - x_1)k_1 \right) \right] + \frac{1}{2}\sigma^2 \Delta W(x). \end{aligned}$$

We abstract from surplus-sharing rules within the cartel and take the cooperative regime as given. Let R denote the reflection in the diagonal, that is, $Rx = R(x_1, x_2) = (x_2, x_1)$. By symmetry of the cooperative problem with respect to firm labels, we have $W(x) = W(Rx)$ and $k_1(x) = k_2(Rx)$. This does clearly not imply that $k_1(x) = k_2(x)$ off the diagonal. At asymmetric states, the cartel may assign different R&D effort levels to the technology leader and the follower; only the firms' labels, not their positions, are irrelevant.

As a special case of the cooperative regime, we consider a research joint venture (RJV). Following the standard approach in the literature (e.g. [Kamien et al., 1992](#); [Samano et al., 2017](#)), we model the RJV as an R&D cartel with full information sharing by setting $\beta^{RJV}(\cdot) \equiv 1$. Firms nevertheless continue to operate separate laboratories, so the RJV corresponds to a secretariat joint venture in the sense of [Martin \(2002\)](#). Under convex R&D costs, a given reduction in marginal cost can be achieved at lower total R&D cost with two laboratories than with one.

3 Numerical Solution Method

We solve the stationary HJB system on a bounded tensor grid in log costs using a semi-Lagrangian discretization ([Falcone, 2003](#); [Falcone and Ferretti, 2013](#)). The scheme is based on the discrete-time dynamic programming principle for controlled diffusions. Over a time increment Δt , we advance the state from each grid point by an Euler step along the control-dependent drift, approximate the conditional expectation induced by the diffusion using a small set of quadrature nodes, and recover off-grid continuation values by linear interpolation. The resulting semi-Lagrangian scheme is monotone and, unlike explicit finite-difference schemes, is not subject to the usual Courant–Friedrichs–Lewy (CFL) time-step restriction, thereby allowing relatively large time steps without compromising stability.

For brevity, we describe the numerical procedure only for the competitive regime. The cartel problem is solved analogously, except that the Bellman update involves joint maximization over (k_1, k_2) . Because the primitives are symmetric across firms, we restrict attention to symmetric

Markov-perfect equilibria, under which

$$V^2(x) = V^1(Rx), \quad k_2(x) = k_1(Rx).$$

This restriction does not impose equal R&D effort at asymmetric states. Rather, it implies that relabeling the firms amounts to interchanging the state variables. It therefore suffices to solve for firm 1's value and policy functions; those for firm 2 follow by symmetry. Accordingly, in what follows, we write, for example, $V(x)$ instead of $V^1(x)$, $\Gamma(x)$ instead of $\Gamma^1(x)$, and $V^n(x)$ instead of $V^{1,n}(x)$.

3.1 Time discretization

The starting point of our numerical scheme is the principle of optimality (14), which can be written as a fixed point equation

$$V = T(V),$$

with the Bellman operator $T(V)$ the right-hand side of (14). This suggests setting up a value-iteration scheme

$$V^{n+1} = T(V^n).$$

That is, we set up the equilibrium as a fixed point of the best-response dynamics: recall that V^n encodes the values of both players.

We perform two approximations to arrive at our numerical method: first we discretize the time step to change the differential game into a continuous-state difference game; in the second step, we discretize the state space.

For the time discretization, we determine Γ^n from V^n and approximate the log-cost dynamics (13) by

$$x(\Delta t, x_0, k, \varepsilon) = x_0 + f(x_0, k, \Gamma^n(Rx_0))\Delta t + \sigma\sqrt{\Delta t} \varepsilon, \quad (16)$$

where $\varepsilon = (\varepsilon_1, \varepsilon_2)$ is a vector-valued random variable whose components are independent and take the values -1 and 1 , each with probability $\frac{1}{2}$. The time-discretized Bellman operator then reads as

$$T_1(V)(x_0) = \max_{k \geq 0} \mathbb{E} \left[(h(x_0) - bk^2)\Delta t + e^{-\rho\Delta t} V(x(\Delta t, x_0, k, \varepsilon)) \right], \quad (17)$$

where $x(\Delta t, x_0, k, \varepsilon)$ is given by (16).

3.2 Space discretization

The second approximation step is to discretize the state variable. We truncate the unbounded log-cost state space to the bounded domain $\Omega = [x_{\min}, x_{\max}]^2$. We discretize Ω using a tensor-product grid with J equispaced nodes in each dimension

$$x_j = (x_{1,j_1}, x_{2,j_2}) = (x_{\min} + (j_1 - 1)\xi, x_{\min} + (j_2 - 1)\xi), \quad j = (j_1, j_2), \quad 1 \leq j_1, j_2 \leq J,$$

where $\xi = (x_{\max} - x_{\min})/(J - 1)$. The approximation of the n th iterate $V^n(x_j)$ of the value at a point x_j is denoted v_j^n . These approximations are collected in the J^2 -dimensional vector $v^n = (v_j^n)_j$.

We face the problem that, in general, even if x_0 is a grid point, the future state $x(\Delta t, x_0, k, \varepsilon)$ will not be a grid point, and the value at this state has to be determined by interpolation. To space-discretize the time-discretized Bellman operator T_1 , we therefore replace the function V on the right-hand side by the bilinear interpolant $I[v]$, which for $x \in [x_{1,j_1}, x_{1,j_1+1}] \times [x_{2,j_2}, x_{2,j_2+1}]$ is given by

$$I[v](x) = \frac{1}{\xi^2} \begin{pmatrix} x_{1,j_1+1} - x_1 & x_1 - x_{1,j_1} \end{pmatrix} \begin{pmatrix} v_{j_1,j_2} & v_{j_1,j_2+1} \\ v_{j_1+1,j_2} & v_{j_1+1,j_2+1} \end{pmatrix} \begin{pmatrix} x_{2,j_2+1} - x_2 \\ x_2 - x_{2,j_2} \end{pmatrix}.$$

Writing h_j for the instantaneous sales profit $h(x_j)$ at grid point x_j , we arrive at the space- and time-discretized Bellman operator

$$\begin{aligned} (T_2(v))_j &= \max_{k \geq 0} \mathbb{E} \left[(h_j - bk^2)\Delta t + e^{-\rho\Delta t} I[v](x(\Delta t, x_j, k, \varepsilon)) \right] \\ &= \max_{k \geq 0} \frac{1}{4} \sum_{m_1=0}^1 \sum_{m_2=0}^1 \left[(h_j - bk^2)\Delta t + e^{-\rho\Delta t} I[v] \left(x(\Delta t, x_j, k, ((-1)^{m_1}, (-1)^{m_2})) \right) \right]. \end{aligned}$$

The maximizer k^* of k is found by discretized search over a finite action set $\mathcal{K} \subset [0, k_{\max}]$. This ensures that the global maximizer is identified even when the objective may not be globally concave. In the implementation, we first perform a coarse global search to locate the neighborhood of the maximizer and then refine locally on a finer control grid. We increase $k_{\max}$ whenever necessary to ensure that the maximum is not attained at the boundary.

Starting from an initial guess, we iterate the Bellman update until the sup-norm difference

$$\|v^{n+1} - v^n\|_{\infty} = \max_j |v_j^{n+1} - v_j^n|$$

falls below a prescribed tolerance, and then recover the equilibrium policy from the maximizing action in the converged solution.⁴

4 Numerical Results

We proceed in two steps. We first isolate the role of state-dependent spillovers by comparing the alternative spillover specifications under noncooperative R&D. We then introduce R&D

⁴All computations are carried out on a 200×200 tensor grid in log-cost space over the domain $[-10, 10]^2$, with a semi-Lagrangian time step of 0.01 and a convergence tolerance of 10^{-8} . The computational domain is chosen sufficiently wide that boundary effects are negligible in the economically relevant region. Refinement checks suggest that this discretization is sufficient for the results reported here and that further refinement does not alter the qualitative conclusions. For the parameterization used in Fig. 2, for instance, reducing the time step from 0.01 to 0.005 yields a mean absolute difference of about 0.001 in the R&D policy function, which is small relative to the typical range of equilibrium R&D levels. Tightening the convergence tolerance from 10^{-8} to 10^{-10} yields a mean absolute difference of only 3.2×10^{-8} .

cooperation and examine how cartelization and RJVs modify the resulting investment incentives, industry dynamics, and welfare.

The figures in the main text are based on the baseline parameterization in Table 1, which we use to illustrate the main mechanisms. The numbered Results below are reserved for qualitative patterns that remain stable across extensive numerical robustness checks. Appendix A reports representative checks under alternative spillover bounds, faster spillover decay, lower R&D costs, and higher discounting and obsolescence.

Table 1: Baseline parameterization

A	b	ρ	δ	σ	β_L	β_H	κ	α
1	6	0.05	0.05	0.2	0.2	0.8	1	0.5

We normalize market size by setting the demand intercept to $A = 1$. We fix $\sigma = 0.2$, which introduces a moderate stochastic component into cost dynamics while leaving the drift, and hence firms' investment decisions, as the main determinant of cost evolution. We set $\rho = \delta = 0.05$, corresponding to moderate annual discounting and knowledge obsolescence rates.

In a continuous-time setting, firms can smooth R&D investment over long horizons; with moderate knowledge depreciation, even small but persistent investment may gradually generate substantial reductions in initially high marginal costs. We therefore set the R&D cost parameter to $b = 6$, which yields a nontrivial region of active investment above the market-entry threshold while making R&D negligible at very high marginal costs. Throughout, state-space figures are presented in marginal-cost coordinates, with logarithmic scaling on both axes.

4.1 Constant spillovers

We begin with the benchmark case of constant spillovers, $\bar{\beta} = 0.2$. This specification provides a reference point for the state-dependent spillover specifications considered below. Because the constant-spillover case is a benchmark rather than the main contribution of the paper, we keep the discussion brief and focus on the equilibrium policy functions and industry dynamics needed to interpret the state-dependent results.

In Fig. 2, firm 1's value is decreasing in its own marginal cost and increasing in the rival's cost. Away from the diagonal, low values of c_1 combined with high values of c_2 place firm 1 in the monopoly region, as indicated by the Cournot entry–exit curves E_1 and E_2 shown in Fig. 3a. The value surface drops sharply when the state crosses the relevant entry boundary and the rival becomes active. The policy surface and contour plot show that equilibrium R&D is highest at relatively high cost levels and then declines smoothly as costs fall. This pattern reflects investment smoothing in continuous time: firms undertake substantial R&D when costs are high and then reduce investment gradually as they become more efficient.

Since equilibrium R&D policies determine the drift of the marginal-cost process in (6), the drift vector field describes the local direction of expected cost evolution at each state. In Fig. 3a, it is shown as velocity vectors $(\dot{c}_1, \dot{c}_2)$ evaluated on a grid in the (c_1, c_2) state space; longer arrows




Figure 2: Equilibrium value and R&D policy under constant spillovers. Panel (a) shows firm 1's value function, panel (b) its equilibrium R&D effort, and panel (c) the corresponding R&D contours, all as functions of marginal costs (c_1, c_2) for $\bar{\beta} = 0.2$.

indicate faster movement. We use the implied drift paths and the associated separatrices to characterize industry dynamics over finite horizons and to identify the regions of the state space associated with the long-run outcomes labelled in Fig. 3a.⁵ This interpretation is also consistent with the small-noise approximation: when the diffusion term is small, solutions of the stochastic system remain close to the corresponding deterministic solutions over finite horizons (Freidlin and Wentzell, 2012, Theorem 1.2).

The solid curves in Fig. 3a are the separatrices of the drift vector field. They divide the state space into four basins of attraction. The basin labelled *Eventual Duopoly* contains the unique stable steady state, shown by the large dot on the diagonal. Trajectories starting in this region therefore converge to long-run coexistence even when firms begin from asymmetric cost positions. Outside this basin, either both firms are too inefficient to develop the technology profitably, or the initial cost gap is too large for the laggard to catch up, leading instead to *No Market* or to *Eventual Monopoly* by one of the two firms. In the latter cases, the firm that gives up on development experiences unbounded marginal-cost growth because its equilibrium R&D effort is insufficient to offset the positive depreciation rate. We interpret this outcome as exit from the market.

Although the separatrices lie close to the Cournot entry and exit curves E_1 and E_2 , they do not coincide with them. The reason is that E_1 and E_2 are static product-market boundaries, whereas the separatrices are dynamic objects determined by equilibrium R&D incentives and the resulting law of motion for costs. As a result, states with similar current product-market status may still lead to different long-run industry structures. In particular, a firm may initially be at a cost disadvantage and still remain in the *Eventual Duopoly* basin if the gap is sufficiently small, so that continued investment allows the laggard to catch up. If the asymmetry is too pronounced, however, the drift points toward eventual monopoly by the more efficient firm.

This dynamic partition of the state space provides the foundation for the analysis that follows.

⁵In Fig. 3a, the $\dot{c}_1 = 0$ and $\dot{c}_2 = 0$ nullclines intersect at four critical points of the drift vector field: two saddles, one nodal source, and one nodal sink (the stable equilibrium), of which only the latter is shown. The separatrices coincide with the relevant nullcline branches through the nodal source and with the stable manifolds of the two saddles.

We now examine how distance-dependent spillovers reshape these basin boundaries and, with them, the industry's long-run dynamics.

4.2 State-dependent spillovers

We now turn to distance-dependent spillovers. We compare the symmetric decay specification in Eq. (4), with $\beta_H = 0.8$, $\beta_L = 0.2$, and $\kappa = 1$, to the asymmetric specification in Eq. (5), evaluated at $\alpha = 0.5$ and $\alpha = 0.25$. In both cases, spillovers are highest when firms are technologically close and decline as the cost gap widens. Lower values of α make spillovers more follower-favoring by reducing the extent to which technological distance shields the leader from imitation. The equilibrium effect is ambiguous a priori, because stronger spillovers to the laggard also change the leader's incentive to invest.

The parameter α can be interpreted as a feature of the industry's knowledge-transmission environment. Lower values describe industries in which lagging firms continue to learn from frontier firms despite being technologically behind, for example because production knowledge is more codified, worker mobility is higher, or process innovations are easier to imitate and adapt.

Figure 3b shows that distance-dependent spillovers shift the basin boundaries in a systematic way. Relative to the symmetric decaying-spillover case, lower values of α enlarge the region of initial states leading to *Eventual Duopoly* and contract the monopoly regions. By contrast, the duopoly steady state itself changes only slightly across specifications, with lower values of α implying a marginally higher long-run cost level. Spillover asymmetry therefore matters mainly through transitional dynamics and market selection rather than through the eventual cost level.⁶

Figure 4 reports differences in firm 1's equilibrium R&D effort across spillover specifications. The economically relevant differences are concentrated above the diagonal, where firm 1 is the technological leader. More follower-favoring spillovers therefore primarily reduce the leader's equilibrium R&D effort in technologically asymmetric states. This is a free-riding effect: as the follower captures a larger share of the benefits from the leader's R&D, the leader's private return to investment falls. The implication for catch-up is ambiguous, however, because lower leader investment also reduces the R&D flow available for imitation. In the computed equilibrium, the increase in the follower's access to spillovers dominates, so lower values of α enlarge the set of initial states from which the industry converges to eventual duopoly.

To show that asymmetric spillover decay affects not only long-run market selection but also the speed of convergence, Fig. 5 reports deterministic transition paths generated by the equilibrium drift. These paths abstract from stochastic shocks and should therefore be interpreted as benchmark trajectories rather than realized sample paths. We consider an initial condition with substantial cost asymmetry, $(c_1, c_2) = (0.8, 1.6)$, that nevertheless lies in the basin of attraction of eventual duopoly under both values of α . This allows a direct comparison of how the degree

⁶Relative to the constant-spillover cases, the symmetric decaying-spillover specifications generate duopoly-basin boundaries that are broadly bracketed by the corresponding boundaries under $\bar{\beta} = 0.8$ and $\bar{\beta} = 0.2$. As κ increases, these boundaries move closer to the low-spillover case $\bar{\beta} = 0.2$. See Appendix A.1, Figure 10, for a more detailed comparison.




Figure 3: Industry dynamics under alternative spillover specifications. Panel (a) shows the drift vector field under constant spillovers with $\bar{\beta} = 0.2$, together with the Cournot entry–exit curves E_1 and E_2 , the stable steady state, and the separatrices dividing the state space into the basins of attraction of no market, eventual monopoly, and eventual duopoly. Panel (b) compares the separatrices under symmetric spillover decay (solid) and asymmetric decay with $\alpha = 0.5$ (dashed) and $\alpha = 0.25$ (dotted). Lower values of α enlarge the basin of attraction of eventual duopoly.

of follower-favoring spillovers shapes catch-up dynamics and the transition toward long-run coexistence. In addition to marginal costs, output, and individual R&D effort, we also report total R&D, $k_1 + k_2$, and total effective R&D, $K_1 + K_2$.

Lowering α increases the follower’s access to spillovers when firms are technologically far apart. This accelerates the follower’s cost reduction, even though it reduces the leader’s private return to R&D and therefore lowers the leader’s own investment. This is not simply because a slower-moving leader is easier to catch. Lower leader investment also reduces the R&D flow available for imitation. The relevant object is therefore the follower’s effective R&D. Panel (d) shows that total R&D is lower under $\alpha = 0.25$ throughout the transition, but total effective R&D is higher in the early phase, when stronger spillovers to the follower more than offset the decline in leader investment. This early increase in effective knowledge transmission accelerates catch-up, brings forward the follower’s entry, and shortens the monopoly phase.

A notable feature of the transition is that the leader initially invests more than the follower, whereas the ranking later reverses. Early in the transition, the leader has the higher private return to cost reduction because it is already producing and can exploit its technological advantage more profitably. During the catch-up phase, additional cost reduction becomes relatively more valuable for the laggard, so the follower’s equilibrium R&D effort eventually exceeds that of the leader. In absolute terms, both firms reduce R&D over time as the cost gap narrows and the industry approaches the steady state. Later in the transition, both total R&D and total effective




Figure 4: Contour plots of differences in firm 1's equilibrium R&D effort across spillover specifications. Panel (a) shows $k_1^{\alpha=0.5} - k_1^{\text{sym}}$. Panel (b) shows $k_1^{\alpha=0.25} - k_1^{\alpha=0.5}$. The visible differences are predominantly negative and concentrated above the diagonal (thin gray line), where firm 1 is the technological leader. Values with absolute magnitude below 10^{-3} are suppressed.

R&D are slightly lower under $\alpha = 0.25$, consistent with its marginally higher long-run cost level.

The output paths mirror this transition. Initially, only the leader produces, so the industry is effectively in a monopoly phase and the leader's output rises as its cost falls. Under lower α , the follower enters production sooner; once this occurs, the market switches to duopoly, the leader's output declines from its pre-entry peak, and the follower's output continues to rise. Over time, both firms converge toward the steady state. Thus, even after firms adjust their equilibrium R&D choices, more follower-favoring spillovers erode market dominance more quickly and make eventual duopoly more likely.

Result 1. *More follower-favoring spillovers (lower α) make eventual duopoly more likely and accelerate technological catch-up from asymmetric initial conditions, despite weakening the leader's incentive to invest.*

The results above show that spillover asymmetry affects industry dynamics primarily by changing how strongly initial technological asymmetries translate into long-run market structure and by altering the speed of technological catch-up. We now turn to R&D cooperation and ask how cartelization modifies these investment incentives, transition dynamics, and market-structure outcomes.

4.3 R&D cooperation

A long-standing result in the R&D cooperation literature is that the effect of cooperation depends critically on the strength of spillovers. In the seminal model of [d'Aspremont and Jacquemin \(1988\)](#), and in subsequent generalizations surveyed by [De Bondt \(1997\)](#), R&D cooperation




Figure 5: Deterministic transition paths under asymmetric spillover decay. The paths are generated by the equilibrium drift from the initial state $(c_1(0), c_2(0)) = (0.8, 1.6)$ for $\alpha = 0.5$ and $\alpha = 0.25$. Panels (a)–(c) report marginal costs, production, and individual R&D effort, respectively. Panel (d) reports total R&D, $k_1 + k_2$, and total effective R&D, $K_1 + K_2$.

tends to reduce investment and industry output when spillovers are low, but to raise them when spillovers are high. At low spillovers, noncooperative firms appropriate most of the gains from their own cost reductions and therefore invest aggressively; cooperation internalizes the effect of these cost reductions on subsequent product-market rivalry. At high spillovers, by contrast, noncooperative firms underinvest because much of the return to innovation spills over to rivals, and cooperation internalizes this knowledge externality. We revisit this logic in our global dynamic setting. In our model, cooperation affects not only the level of R&D, but also the evolution of relative costs and hence the long-run market structure induced by the cost dynamics.

4.3.1 Low constant spillovers

For low spillovers, our dynamic model confirms the standard concern that R&D cooperation need not improve market outcomes. For $\bar{\beta} = 0.2$, cartel R&D falls below noncooperative R&D over much of the state space, particularly in the low-cost duopoly region (Fig. 7a). More consequentially, cartelization changes long-run market selection. Whereas the noncooperative




Figure 6: Drift vector fields under the R&D cartel. Panel (a) shows the cartel with low constant spillovers, $\bar{\beta} = 0.2$. Panel (b) shows the RJV cartel with full information sharing, $\beta^{RJV} = 1$; dotted curves indicate the corresponding RJV competition loci.

benchmark in Fig. 3a yields a wide basin of attraction around the stable low-cost duopoly steady state, the cartel shrinks this basin to the diagonal itself. In Fig. 6a, the low-cost diagonal steady state, indicated by the gray dot, is attracting only along the diagonal: off-diagonal perturbations are repelled, so small cost asymmetries are amplified and the industry is steered toward monopoly by one firm.

The mechanism is straightforward. Product-market competition remains Cournot, so joint profits can be higher when one sufficiently efficient firm serves the market than when both firms remain active. The cartel therefore faces two opposing forces. Convex R&D costs make it efficient to use both laboratories, but the joint-profit objective also gives the cartel an incentive to move the industry toward states in which production is concentrated in the lower-cost firm. When spillovers are low, the second force dominates: coordinated R&D reinforces the lower-cost firm rather than preserving cost parity. R&D cooperation can therefore induce monopolization even without explicit product-market collusion.

This effect is also visible on the diagonal. Relative to the noncooperative benchmark in Fig. 3a, the low-spillover cartel has a higher low-cost diagonal steady state, consistent with the broad reduction in R&D effort shown in Fig. 7a.

Result 2. *At low spillovers, R&D cartelization reduces investment over much of the duopoly region, amplifies cost asymmetries, and steers the industry toward monopoly by one firm.*

4.3.2 High constant spillovers and the RJV limit

As spillovers increase, R&D performed in either laboratory becomes more valuable for joint cost reduction. This matters especially in high-cost regions, where production is not yet viable or one




Figure 7: Contour plots of the difference between cartel and noncooperative R&D effort under constant spillovers. The panels show $k_1^{\text{Cartel}}(x_1, x_2) - k_1^{\text{Comp}}(x_1, x_2)$. Panel (a) sets $\bar{\beta} = 0.2$, while panel (b) sets $\bar{\beta} = 0.8$. White regions indicate absolute differences below 10^{-3} . The solid line is the diagonal, and the dashed curves are the Cournot entry–exit curves E_1 and E_2 .

firm is close to entering the product market. As shown in Fig. 7b, the cartel invests more than noncooperative firms in these regions, with the increase concentrated on the lagging-firm side of the state space. The laggard’s R&D lowers its own cost and, through spillovers, also lowers the leader’s cost. Coordinated R&D can therefore speed up cost reduction where it is most valuable: when production is within reach or when the leader already produces and lower production costs raise joint profits. This effect is visible in the drift vector field. In our computations, R&D cooperation consistently reduces the no-market region relative to noncooperative R&D, and the reduction is larger when spillovers are higher. Thus, even when cartelization weakens future product-market rivalry, it makes technologies with high initial costs more likely to be developed.

Result 3. *R&D cartelization makes high-cost technologies more likely to be brought to market by reducing the no-market region relative to noncooperative R&D. This effect is stronger when spillovers are higher.*

High spillovers also slow the transition toward monopoly. The corresponding drift field, omitted to save space, has the same qualitative geometry as Fig. 6a, but with weaker movement away from the diagonal. When $\bar{\beta}$ is high, drift vectors close to the diagonal are almost parallel to it, so trajectories that start with small cost asymmetries can remain in the Cournot duopoly region for a long time. But this persistence is not permanent unless spillovers are complete. As long as $\bar{\beta} < 1$, relative-cost drift remains: once one firm gains a cost advantage, the cartel eventually reinforces that advantage rather than offsetting it. The reason is that, at low cost levels, the common-cost benefit from spreading R&D across both laboratories is weaker. The cartel then has less incentive than noncooperative firms to preserve a duopoly. Instead, further reducing the leader’s cost can raise joint profits by moving the industry toward a more concentrated

market structure. This is reflected in the negative cartel–competition R&D differences below the diagonal in Fig. 7.

Higher spillovers make the cartel more effective at generating common cost reductions, but they do not eliminate its tendency to amplify cost asymmetries later in the development process. The steady-state comparison illustrates the cost effect. With low spillovers, the cartel’s low-cost diagonal steady state lies above the corresponding noncooperative steady state. With high spillovers, the ranking reverses: the cartel moves the industry toward lower diagonal costs. Yet incomplete spillovers still leave a residual force toward asymmetry. Higher spillovers weaken this force and delay the transition toward monopoly; low spillovers strengthen it and accelerate monopoly selection.

Result 4. *Within the R&D cartel, high spillovers lower costs while firms remain technologically close and delay the transition toward monopoly. Low spillovers imply higher costs near parity and accelerate that transition.*

The RJV cartel in Fig. 6b is the limiting case. Full information sharing sets $\beta^{RJV} = 1$, which equalizes the firms’ deterministic log-cost drifts, $\dot{x}_1 = \dot{x}_2$. The log-cost gap $x_1 - x_2$ is therefore invariant, so the deterministic trajectories appear parallel to the diagonal on the log-scaled cost axes. Since R&D performed in either laboratory has the same effect on both firms’ costs, the computed firm-level policy exhibits the expected mirror symmetry, $k_1(x_1, x_2) = k_1(x_2, x_1)$. The two zero-drift nullclines, $\dot{x}_1 = 0$ and $\dot{x}_2 = 0$, therefore coincide, and the low-cost steady-state outcome becomes a one-dimensional equilibrium set rather than an isolated steady state. This equilibrium set is shown by the blue locus running from the upper-left to the lower-right of Fig. 6b. Along this set, the Jacobian has one negative eigenvalue in the transverse direction and one zero eigenvalue in the tangent direction. The equilibrium set is therefore attracting but non-hyperbolic: trajectories are drawn toward it, while the initial technology gap selects the long-run cost configuration.

This geometry differs from the incomplete-spillover case. Full information sharing removes the relative-cost drift that, for $\bar{\beta} < 1$, may be slow but still eventually pushes asymmetric states toward monopoly. The RJV cartel therefore generates a genuine eventual-duopoly region: for a range of asymmetric initial cost positions, the cost gap remains bounded and both firms remain active in the long run.

Relative to RJV competition, defined as noncooperative R&D with $\beta = 1$, the RJV cartel dominates on all dimensions shown in Fig. 6b. The dotted loci show the basin boundaries and equilibrium set under RJV competition. The cartel generates a larger eventual-duopoly region, a smaller no-market region, and lower long-run costs along the attracting equilibrium set.

Result 5. *Among the R&D cartel regimes considered, the RJV cartel is most effective at developing high-cost technologies: it generates the smallest no-market region and leads to the lowest long-run costs. It also yields the largest eventual-duopoly basin.*

This dominance does not extend to all noncooperative regimes. At high cost levels, the RJV cartel generates the widest eventual-duopoly basin, because full information sharing and the coordinated use of both laboratories make development most effective. At low cost levels,

however, comparing Fig. 6b with the low-spillover competitive benchmark in Fig. 3a shows that noncooperative R&D with low spillovers can generate a wider eventual-duopoly basin. The reason is that full information sharing makes the log-cost gap invariant: the RJV cartel lowers common costs efficiently, but it does not generate relative catch-up. Low-spillover competition, by contrast, can induce stronger catch-up incentives for the laggard in some asymmetric states. Thus, RJV cooperation improves joint technology development and lowers long-run costs, but it need not uniformly enlarge the set of initial states leading to eventual duopoly.

Result 6. *The RJV cartel does not uniformly dominate noncooperative R&D in preserving duopoly. At lower cost levels, noncooperative R&D can generate stronger catch-up incentives and a larger eventual-duopoly basin.*

4.3.3 Follower-favoring spillovers

Distance-decaying spillovers introduce an additional mechanism for preserving duopoly under the R&D cartel. The logic is related to the noncooperative case, but it now works against the cartel's tendency to reinforce the current leader. Near the diagonal, where firms have similar costs, spillovers are high and R&D in either laboratory generates large common cost reductions. Follower-favoring asymmetry strengthens this stabilizing effect because the laggard receives more of the leader's R&D than the leader receives of the laggard's R&D. For sufficiently small gaps, trajectories remain in the duopoly region rather than drifting toward monopoly.

Figure 8a illustrates the cartel dynamics for $\alpha = 0.5$, together with the corresponding noncooperative separatrices, shown as dotted curves. As in the constant-spillover case, R&D cooperation reduces the no-market region, making development of initially high-cost technologies more likely. The cartel also converges to a lower diagonal steady state than the noncooperative regime, reflecting stronger incentives to generate common cost reductions when spillovers are high near the diagonal.

At the same time, cartelization does not uniformly expand the duopoly region. Away from the high-cost development margin, the cartel still internalizes product-market rivalry and therefore has weaker incentives to maintain two active firms. As a result, the cartel's eventual-duopoly basin remains smaller than the noncooperative one over much of the state space, with the gap between the two basin boundaries widening as costs fall.

Figure 8b shows how the cartel basin changes as spillovers become more follower-favoring. Reducing α from 0.5 to 0.25 moves the duopoly-basin boundaries farther from the diagonal, while leaving the diagonal steady state nearly unchanged. Stronger spillovers to the laggard offset the cartel's tendency to reinforce the current leader, expanding the set of asymmetric states from which the industry converges to duopoly.

The RJV cartel again provides the natural benchmark. Full information sharing eliminates relative-cost drift and, in our numerical analysis, yields the smallest no-market region, the largest eventual-duopoly basin, and the lowest long-run costs among the cartel regimes considered (Appendix A.3, Figure 13, reports the same ranking under alternative specifications). The distance-decay cases show, however, that complete information sharing is not necessary for




Figure 8: Follower-favoring spillovers under the R&D cartel. Panel (a) shows the cartel drift vector field for $\alpha = 0.5$, with the corresponding competition separatrices shown as dotted curves. Panel (b) compares cartel separatrices for $\alpha = 0.5$ (solid) and $\alpha = 0.25$ (dotted). Lower values of α expand the basin of attraction of eventual duopoly.

R&D cooperation to preserve duopoly from asymmetric initial conditions. Follower-favoring asymmetry strengthens this duopoly-preserving effect: stronger asymmetry in spillover decay enlarges the local basin around the diagonal steady state and makes duopoly robust to larger technology gaps.

Result 7. *Under distance-decaying spillovers, the R&D cartel can preserve duopoly locally: for sufficiently small technology gaps, cost trajectories remain in the duopoly region rather than drifting toward monopoly. Stronger follower-favoring asymmetry expands the eventual-duopoly basin.*

The analysis therefore leaves a welfare question open. R&D cooperation can make development more likely and lower long-run costs, but it can also weaken the forces that preserve duopoly. Whether this tradeoff is desirable depends on its implications for consumer surplus, producer surplus, and total surplus. We therefore turn to welfare comparisons across regimes.

4.4 Welfare comparisons across R&D regimes

We now compare R&D cooperation with noncooperative R&D under the same knowledge-transmission environment. Welfare is measured using expected present values of consumer surplus, producer surplus, and total surplus. For each initial state, we simulate 1,000 stochastic trajectories of the marginal-cost process using the Euler–Maruyama scheme.⁷ Along each trajectory, instantaneous Cournot consumer surplus is $CS_t = \frac{1}{2}Q_t^2$, where $Q_t = q_1(t) + q_2(t)$

⁷To improve the precision of cross-regime comparisons, we use common random numbers: for a given initial state, all regimes are simulated using the same Brownian shock realizations.

is total output. Discounted consumer surplus is computed by applying the trapezoidal rule to each simulated path and then averaging across trajectories. Expected producer surplus (profit) is obtained directly from the converged value functions evaluated at the initial state. Total surplus is the sum of expected consumer surplus and expected producer surplus.

Because the model has an infinite horizon, simulations are truncated once the discounted surplus increment over a time step falls below 10^{-6} . If a simulated trajectory reaches the boundary of the computational domain, instantaneous surplus is held constant thereafter.

To assess how initial asymmetry shapes welfare, we compare outcomes for initial cost pairs on two iso-sum loci in (c_1, c_2) space, varying $\Delta c \equiv c_1 - c_2$ while holding $c_1 + c_2$ fixed. This keeps the overall initial cost level constant and isolates the effect of asymmetry in initial cost positions. The first locus, ISL_1 , passes through the symmetric state $c_1 = c_2 = 1.3$, where production is initially not viable and R&D must precede entry. The second, ISL_2 , passes through the symmetric state $c_1 = c_2 = 0.4$, where firms begin in an active duopoly. The two loci therefore contrast welfare effects for high-cost technologies with those for already viable technologies. By symmetry, we report only the positive half of each locus, varying the initial cost difference Δc .

The two iso-sum loci are chosen so that all regimes develop the technology under the deterministic drift classification. The welfare comparisons therefore focus on how R&D cooperation changes surplus where development is possible under all regimes, rather than on the discrete gain from bringing to market a technology that another regime would leave undeveloped. As the drift analysis showed, regimes with smaller no-market regions can develop technologies from higher initial costs. For sufficiently high initial costs, such regimes have a direct welfare advantage because they generate surplus in states where other regimes do not.

Figures 9a and 9b report cartel–competition welfare differences along ISL_1 and ISL_2 , respectively, for each spillover specification. Positive values indicate that R&D cooperation raises the corresponding surplus measure relative to the noncooperative benchmark. In our computations, cartelization raises producer surplus at every initial state on the two iso-sum loci. The central question is therefore whether these private gains reflect dynamic efficiency improvements or instead come at the expense of consumers.

The answer depends strongly on the initial state and on the spillover environment. With low constant spillovers, the model confirms the standard concern that R&D cooperation can benefit firms without benefiting consumers. For $\bar{\beta} = 0.2$, consumer surplus falls along the plotted loci, and total surplus is lower in most of the plotted states, although the loss becomes small and can turn into a gain at sufficiently asymmetric states. Appendix A.4 reports alternative parameterizations showing that the precise location of this sign change varies with the calibration, but the underlying pattern is robust: as asymmetry increases, the consumer-surplus penalty from cartelization weakens.

The mechanism is that, at low spillovers, the cartel has little knowledge externality to internalize. Near symmetry, noncooperative firms have strong private and business-stealing incentives to invest, which benefit consumers by accelerating cost reduction and preserving future rivalry. By




Figure 9: Welfare effects of R&D cartelization. Positive values indicate gains from R&D cooperation relative to the corresponding noncooperative regime.

weakening these incentives, the cartel tends to reduce R&D effort and effective support for the high-cost firm, allowing cost asymmetries to widen and reducing consumer surplus. As the initial cost gap increases, however, the high-cost firm is already a weaker competitive threat, especially when knowledge flows are limited. The noncooperative benchmark therefore loses part of its advantage, so producer-surplus gains can dominate in some sufficiently asymmetric states. Overall, low constant spillovers provide a weak case for R&D cooperation: with little knowledge externality to internalize, producer-surplus gains are often offset by consumer-surplus losses.

The pattern reverses when spillovers are high. For $\bar{\beta} = 0.8$, cartelization raises consumer surplus and total surplus over most of the plotted domain, with losses confined to highly asymmetric states in some specifications. In this case, R&D cooperation has a clear efficiency rationale. At high spillovers, the cartel internalizes large knowledge externalities, raising joint and effective R&D over much of the relevant state space. The induced cost dynamics move the industry toward lower costs while keeping the laggard technologically relevant over a wider range of

initial states. These gains decline with the initial cost gap. As the laggard falls farther behind, additional coordinated R&D is less likely to sustain production by both firms, and the cartel's R&D advantage over the noncooperative regime becomes smaller. The consumer-surplus gain from preserving future rivalry therefore weakens.

Result 8. *The welfare effect of R&D cartelization depends on the spillover environment and on the initial state. At low spillovers, cartelization benefits firms but reduces consumer surplus and usually total surplus. At high spillovers, cartelization has a clear efficiency rationale: by internalizing dynamic knowledge spillovers, it raises consumer surplus and total surplus for most initial conditions considered, although the gains tend to decline with initial cost asymmetry.*

The distance-dependent cases are intermediate. Near the diagonal, spillovers are high and reciprocal, so R&D performed in either laboratory generates substantial common cost reductions. Noncooperative firms benefit from these spillovers but do not fully internalize the common-cost-reduction benefit. Cartelization therefore has an efficiency role near symmetry. In the policy functions, it raises joint effective R&D; at the diagonal and over the relevant near-symmetric region, the drift field shows faster downward movement in costs rather than movement away from duopoly. Because both firms remain plausible future competitors, these common cost reductions can raise consumer surplus as well as producer surplus and total surplus.

Follower-favoring decay strengthens this mechanism. Relative to symmetric decay, the $\alpha = 0.5$ cartel generates larger surplus gains because knowledge flows more strongly toward the laggard, making it easier to keep the high-cost firm technologically relevant and preserve effective duopoly.⁸ This advantage is state dependent. As the initial cost gap widens, spillovers decay and the high-cost firm becomes harder to keep technologically relevant. The problem is especially acute at lower cost levels, where, as shown in the previous section, the gap between the noncooperative and cartel duopoly boundaries widens. In this region, preserving the laggard requires R&D that mainly serves to sustain future product-market rivalry, making it less attractive from the cartel's joint-profit perspective. The policy and drift comparisons reflect this shift: cartelization no longer reliably raises joint or effective R&D relative to competition, and the induced dynamics are less likely to move the industry toward low-cost duopoly. As a result, consumer-surplus and total-surplus gains may turn negative, even though firms still gain from cartelization.

Result 9. *With distance-dependent spillovers, R&D cartelization is most beneficial near symmetry, where knowledge flows are strong and R&D generates large common cost reductions. Follower-favoring spillovers increase these gains by directing knowledge toward the laggard and making future duopoly more likely. As the initial technology gap widens, especially when production is already viable, the gains decline and may turn negative because spillovers weaken and the cartel has less incentive to keep the laggard technologically relevant.*

⁸Additional computations with stronger follower-favoring spillovers, $\alpha = 0.25$, reinforce this mechanism: relative to the plotted $\alpha = 0.5$ case, stronger follower-favoring spillovers raise the cartelization gains in consumer surplus, producer surplus, and total surplus along both iso-sum loci. Faster spillover decay, $\kappa = 3$, works in the opposite direction: it attenuates the gains from distance-dependent spillovers, so consumer-surplus and total-surplus gains decline more quickly with the initial cost gap and may become negative, especially along ISL_2 .

The RJV provides the strongest welfare case for R&D cooperation among the regimes considered. It generates the largest average increases in consumer surplus and total surplus, and this ranking is the most robust across the parameterizations considered. Producer surplus also rises, although the incremental producer-surplus gain is not uniformly the largest because each cartel is evaluated relative to its own noncooperative benchmark. The welfare gains are strongest near symmetric states, where both laboratories remain productive and full information sharing allows each firm's R&D to reduce both firms' costs. In this environment, cartelization raises investment in a shared R&D technology and directs the induced cost dynamics toward lower common costs, rather than primarily softening future product-market rivalry. As a result, coordinated R&D leads more consistently to higher output and higher welfare, and its benefits persist over the widest range of asymmetries. Thus, among the cooperation regimes considered, the RJV offers the clearest dynamic-efficiency rationale for coordination.⁹

Result 10. *Among the cooperation regimes considered, the RJV provides the strongest welfare case for coordination. Full information sharing makes R&D benefits common across firms and delivers the most robust gains in consumer surplus and total surplus. It therefore makes cooperation more likely to improve dynamic efficiency, rather than mainly softening rivalry and shifting surplus toward firms.*

The policy message is therefore conditional on the spillover environment rather than uniform. R&D cooperation is most likely to improve welfare when it creates substantial knowledge transmission, rather than merely coordinating firms' R&D choices. This condition is clearest in the RJV benchmark, where full information sharing makes R&D benefits common across firms and gives cooperation its strongest welfare rationale. By contrast, cooperation with limited knowledge exchange offers a weaker welfare case. With little knowledge externality to internalize, its main effect is often to reduce the high noncooperative R&D effort generated by competitive pressure. The private profitability of cooperation is then a poor guide to its social value.

The state-dependent spillover results sharpen this message. The welfare effect of R&D cooperation depends not only on the amount of knowledge transmitted, but also on its direction and persistence as firms move apart. Follower-favoring spillovers improve the performance of cooperation by channeling knowledge toward the firm whose continued technological relevance sustains future duopoly. This mechanism is fragile: even when spillovers are strong near symmetry, stochastic cost improvements by one firm can widen the technology gap, weaken subsequent knowledge transmission, and reduce the cartel's incentive to preserve the laggard. The fragility is most costly at low-cost asymmetric states, where production is already viable and the laggard's main

⁹This ranking echoes the classic RJV rationale in [Kamien et al. \(1992\)](#) and the dynamic results of [Samano et al. \(2017\)](#), who study constant spillovers in a discrete-time model and find that welfare gains from cooperation increase with spillovers and are largest under full information sharing. The present global dynamic framework adds three channels. First, because R&D can precede production, cooperation affects whether high-cost technologies are developed; in our computations, the RJV is also the regime most likely to bring such technologies to market. Second, cooperation affects endogenous market structure: the industry may converge to duopoly or to production concentrated in one firm. Third, the cartel chooses state-contingent firm-level R&D policies, so support for the laggard need not equal support for the leader once a technology gap opens. With low spillovers, the cartel has little knowledge externality to internalize and weaker incentives to support the laggard, making production more likely to concentrate in the technological leader. Low-spillover welfare effects are therefore substantially less favorable than in settings where cooperation changes R&D effort without altering product-market participation.

welfare role is to preserve future product-market competition.

This is why constant-spillover benchmarks can be misleading for welfare comparisons. A low constant-spillover specification can understate the efficiency role of cooperation near symmetry, where technologically close firms tend to exchange knowledge more intensively and coordinated R&D generates large shared cost reductions. A high constant-spillover specification can instead overstate the robustness of cooperation gains, because it assumes that knowledge continues to flow even after firms become technologically distant. When those knowledge flows weaken, the cartel has less incentive to support the laggard, and cooperation is more likely to reinforce concentration in the leading firm. The welfare case for R&D cooperation therefore depends not only on the overall strength of spillovers, but also on how spillovers decay with technological distance and whether they continue to support lagging firms whose presence helps preserve future competition.

5 Conclusion

This paper shows why a fixed spillover parameter alone is not enough to assess the welfare effects of R&D cooperation. We develop a global dynamic model of duopolistic process innovation in which spillovers depend on firms' technological positions. The model nests the standard constant-spillover benchmark but also allows knowledge transmission to vary with technological distance and to differ between leaders and followers. This state dependence creates a feedback loop: R&D changes the cost gap, the cost gap changes spillover intensity, and spillover intensity changes subsequent R&D incentives. We use this framework to compare noncooperative R&D, R&D cartelization, and research joint ventures, and to study how alternative spillover specifications affect catch-up dynamics, technology development, long-run market structure, and welfare.

The main lesson is that the welfare effects of R&D cooperation depend on the level and direction of spillovers. When spillovers are low, there is little knowledge externality to internalize. Cooperation then mainly softens dynamic rivalry: R&D falls over much of the duopoly region, cost asymmetries widen, and the industry is more likely to move toward production by a single firm. When spillovers are high, cooperation has a stronger efficiency rationale. It internalizes knowledge flows, reduces the set of technologies that are never brought to market, lowers costs especially when firms remain technologically close, and can raise consumer surplus and total surplus for a wide range of initial conditions. Thus, cooperation is most likely to improve welfare when it generates substantial information sharing, rather than merely coordinating R&D choices.

State-dependent spillovers sharpen this lesson. More follower-favoring spillovers reduce the leader's private incentive to invest, but they also direct knowledge toward the firm for which it is most valuable. In the computed equilibria, this second force dominates: stronger follower-favoring spillovers accelerate catch-up, shorten monopoly phases, and enlarge the basin of attraction of eventual duopoly. Cooperation is most valuable when firms are technologically close: knowledge flows are then strongest, and coordinated R&D generates the largest common

cost reductions. Follower-favoring spillovers can therefore strengthen the welfare gains from cooperation by helping firms remain close enough for knowledge flows to stay large and for both firms to remain viable future competitors. These gains are fragile, however: as the technology gap widens, spillovers decay, the laggard becomes harder to keep technologically relevant, and cooperation may reinforce concentration rather than preserve rivalry.

The research-joint-venture regime highlights the role of information sharing most clearly. Full information sharing makes the benefits of R&D common across firms. Among the cooperative regimes considered, it yields the smallest no-market region, the lowest long-run costs, and the most robust welfare gains. Yet even this regime does not dominate noncooperative R&D in every respect. At lower cost levels and for sufficiently asymmetric initial conditions, noncooperative R&D may preserve duopoly over a larger set of initial states. The policy implication is that R&D cooperation should be evaluated not only by the coordination it permits, but also by the knowledge flows it creates, their direction, and their persistence as firms move apart.

The two-firm structure is deliberately simple. It allows us to characterize global dynamics and to distinguish no-market, one-firm, and two-firm outcomes. In richer oligopolies, the analogue of monopoly in the present model need not be literal monopoly; it may instead correspond to concentration of production and innovative activity among a small set of technological leaders. R&D cooperation in such settings would also separate knowledge flows inside research coalitions from those between coalition members and outsiders, a distinction likely to matter for innovation incentives and long-run market structure. Extending the analysis to richer oligopoly environments is therefore a natural next step.

Code Availability

Computer code sufficient to reproduce the numerical results in this paper is available at Mendeley Data, <https://doi.org/10.17632/gxhd2jt39r.1>.

Appendix A Robustness checks

This appendix reports robustness checks under alternative parameterizations that vary the main economic forces in the model: spillover decay, R&D costs, discounting, and knowledge obsolescence. The main conclusions are robust across these checks.

A.1 Alternative spillover specifications

Figure 10a compares the separatrices under constant spillovers, $\bar{\beta} = 0.2$ and $\bar{\beta} = 0.8$, with those under symmetric distance-decaying spillovers for $\kappa = 1$ and $\kappa = 3$. In line with [Smrkolj and Wagener \(2019\)](#), raising constant spillovers from $\bar{\beta} = 0.2$ to $\bar{\beta} = 0.8$ enlarges the eventual-duopoly region for high initial cost levels, but not uniformly across the state space, as reflected in the crossing of the corresponding separatrices. At low cost levels, stronger spillovers may significantly weaken the leader's incentive to invest, so that this pro-competitive effect need




Figure 10: Robustness of basin boundaries under faster spillover decay. Panel (a) compares the separatrices under constant spillovers $\bar{\beta} = 0.2$ (solid) and $\bar{\beta} = 0.8$ (dotted) with the symmetric distance-decay specification for $\kappa = 1$ (dash-dotted) and $\kappa = 3$ (dashed), holding $\beta_H = 0.8$ and $\beta_L = 0.2$ fixed. Increasing κ shifts the symmetric-decay separatrices toward the low-spillover case $\bar{\beta} = 0.2$. Panel (b) compares the symmetric decay case (solid) with the asymmetric decay cases $\alpha = 0.5$ (dashed) and $\alpha = 0.25$ (dotted) for $\kappa = 3$. Even with faster spillover decay, lower values of α continue to enlarge the basin of attraction of eventual duopoly relative to the symmetric benchmark.

not persist. The symmetric-decay cases are broadly bracketed by the two constant-spillover cases and move closer to the low-spillover case $\bar{\beta} = 0.2$ as κ increases. Thus, faster decay of spillovers with technological distance reduces the basin of attraction of eventual duopoly under symmetric decay. The corresponding steady states, shown in the figure, also lie between the two constant-spillover benchmarks, with the highest long-run cost under $\bar{\beta} = 0.8$ and the lowest under $\bar{\beta} = 0.2$. This confirms that the effect of spillovers on long-run market structure need not align with their effect on long-run cost reduction among active firms.

Figure 10b considers the asymmetric distance-decay specification with faster overall spillover decay, $\kappa = 3$. Even in this case, more follower-favoring spillovers, corresponding to lower values of α , continue to enlarge the basin of attraction of eventual duopoly relative to the symmetric benchmark. The effect is weaker than in the main text, however, because faster overall decay limits the extent to which the follower benefits from distance-dependent spillovers when firms are far apart.

Figure 11a lowers the spillover bounds to $\beta_H = 0.6$ and $\beta_L = 0.1$, relative to the baseline values $\beta_H = 0.8$ and $\beta_L = 0.2$. Figure 11b lowers the R&D cost parameter from $b = 6$ to $b = 4$, holding the spillover parameters at their baseline values. Lower R&D cost shifts the boundary of the *No Market* region outward, since firms can profitably develop technologies from higher initial cost levels. In both cases, however, the qualitative role of spillover asymmetry is unchanged: lower values of α continue to enlarge the basin of attraction of eventual duopoly relative to the




Figure 11: Robustness under alternative parameterizations. Panel (a) lowers the spillover bounds to $\beta_H = 0.6$ and $\beta_L = 0.1$ and compares the symmetric case (solid) with the asymmetric cases $\alpha = 0.5$ (dashed) and $\alpha = 0.25$ (dotted). Panel (b) lowers the R&D cost parameter from $b = 6$ to $b = 4$, holding the spillover parameters at their baseline values, and compares the same three spillover specifications. In both panels, lower values of α continue to enlarge the basin of attraction of eventual duopoly relative to the symmetric benchmark.

symmetric benchmark.

A.2 Higher discounting and faster obsolescence

We next consider an environment in which long-horizon R&D incentives are weaker because firms are more impatient and knowledge depreciates faster. Specifically, we raise both ρ and δ from 0.05 in the baseline parameterization to 0.1. Under the baseline R&D cost parameter, this removes the interior development region. To retain a nontrivial investment region and keep the comparison informative, we therefore reduce the R&D cost parameter from $b = 6$ to $b = 2$. We then consider both $\kappa = 1$ and $\kappa = 3$. Figure 12 shows that, in both cases, lower values of α continue to enlarge the basin of attraction of eventual duopoly relative to the symmetric benchmark. Faster spillover decay at $\kappa = 3$ attenuates, but does not overturn, this effect.

A.3 Robustness of R&D cartel dynamics

We also check whether the qualitative conclusions for R&D cooperation are robust to alternative parameterizations. Figure 13a considers the higher-discounting, faster-obsolescence case with $\rho = \delta = 0.1$ and, as in Fig. 12, sets $b = 2$ to retain a nontrivial investment region. Figure 13b lowers the R&D cost parameter from $b = 6$ to $b = 4$, holding the remaining baseline parameters fixed. In both panels, the ranking of basin boundaries is unchanged. Lower values of α enlarge the cartel basin of attraction of eventual duopoly, and the RJV cartel generates the widest eventual-duopoly basin and the smallest no-market region. The long-run symmetric duopoly




Figure 12: Robustness under higher discounting and faster obsolescence. Both panels set $\rho = \delta = 0.1$ and $b = 2$. In each panel, the solid curve denotes the symmetric-decay case, the dashed curve the asymmetric case with $\alpha = 0.5$, and the dotted curve the asymmetric case with $\alpha = 0.25$. Lower values of α continue to enlarge the basin of attraction of eventual duopoly relative to the symmetric benchmark.

steady states under $\alpha = 0.5$ and $\alpha = 0.25$ are nearly identical, whereas the RJV equilibrium set lies at lower long-run costs.

A.4 Robustness of welfare results

This subsection reports welfare comparisons under two alternative parameterizations that alter dynamic R&D incentives. The first lowers the R&D cost parameter from $b = 6$ to $b = 4$. The second raises discounting and knowledge depreciation to $\rho = \delta = 0.1$, with $b = 2$ as in Section A.2. In each case, the figures report cartel–competition differences along the same two iso-sum loci used in the main text.

Figure 14 shows that lowering the R&D cost parameter changes the magnitude of the welfare effects but leaves the main consumer-surplus and total-surplus patterns intact. Low-spillover cartelization remains the weakest case for consumers and total surplus on average and over most plotted states, while high spillovers and the RJV generate the most robust total-surplus gains. As in the baseline, the consumer-surplus and total-surplus gains under high spillovers, distance-dependent spillovers, and the RJV tend to decline as the initial cost gap widens. The distance-dependent cases still perform relatively well near symmetry on ISL_1 , but deteriorate on ISL_2 , where production is already viable. With cheaper R&D, noncooperative firms invest more aggressively to preserve future competitive positions, while the cartel has less incentive to support the laggard once spillovers decay with distance.

Figure 15 reports the same welfare comparisons under higher discounting and faster obsolescence. The large positive gains from high spillovers and the RJV are smaller, as future cost reductions






Figure 13: Robustness of R&D cartel basin boundaries. Panel (a) raises the discount rate and knowledge obsolescence rate to $\rho = \delta = 0.1$ and sets $b = 2$. Panel (b) lowers the R&D cost parameter from $b = 6$ to $b = 4$, holding the remaining baseline parameters fixed. Each panel compares follower-favouring spillovers with $\alpha = 0.5$ (dashed) and $\alpha = 0.25$ (dotted) with the RJV cartel (solid). The gray curve denotes the RJV equilibrium set; dots denote the symmetric duopoly steady states under the asymmetric-spillover cartel regimes.

matter less. The main ordering of consumer-surplus and total-surplus gains is nevertheless preserved: the RJV delivers the largest gains, high constant spillovers follow, distance-dependent spillovers remain intermediate, and low constant spillovers provide the weakest case for R&D cooperation.

Across the two alternative parameterizations in Figs. 14 and 15, the producer-surplus component is less uniformly ordered than consumer surplus or total surplus because each welfare difference is measured relative to its own noncooperative benchmark. This is most visible in the lower-R&D-cost calibration on ISL_2 , where low-spillover cartelization can generate larger incremental producer-surplus gains than regimes with stronger knowledge transmission. With low initial costs and cheaper R&D, cartelization under low spillovers mainly softens R&D and product-market rivalry, which is valuable to firms. In regimes with stronger knowledge transmission, the noncooperative benchmark already captures more of the dynamic gains from R&D, so the additional producer-surplus gain from cartelization can be smaller. This does not overturn the main total-surplus ranking. Low-spillover cartelization can generate larger private gains for firms, but these gains are offset by consumer-surplus losses. This reinforces the distinction between private and social gains from R&D cooperation.




(a) ISL_1 : high initial costs




(b) ISL_2 : low initial costs

Figure 14: Welfare effects of R&D cartelization under lower R&D costs. The figure reports cartel-competition differences along ISL_1 and ISL_2 within each spillover specification, setting $b = 4$ and keeping the remaining parameters at their baseline values. Positive values indicate gains from cartelization.



(a) ISL_1 : high initial costs



(b) ISL_2 : low initial costs

Figure 15: Welfare effects of R&D cartelization under higher discounting and faster obsolescence. The figure reports cartel–competition differences along ISL_1 and ISL_2 within each spillover specification, setting $\rho = \delta = 0.1$ and $b = 2$. Positive values indicate gains from cartelization.

References

- ACEMOGLU, D., P. AGHION, AND F. ZILIBOTTI (2006): “Distance to Frontier, Selection, and Economic Growth,” *Journal of the European Economic Association*, 4, 37–74.
- ACS, Z. J., D. B. AUDRETSCH, AND M. P. FELDMAN (1994): “R&D Spillovers and Recipient Firm Size,” *Review of Economics and Statistics*, 76, 336–340.
- ALCÁ CER, J. AND W. CHUNG (2007): “Location Strategies and Knowledge Spillovers,” *Management Science*, 53, 760–776.
- AMIR, R., I. EVSTIGNEEV, AND J. WOODERS (2003): “Noncooperative versus Cooperative R&D with Endogenous Spillover Rates,” *Games and Economic Behavior*, 42, 183–207.
- AMIR, R. AND J. WOODERS (1998): “Cooperation vs. Competition in R&D: The Role of Stability of Equilibrium,” *Journal of Economics*, 67, 63–73.
- BLOOM, N., M. SCHANKERMAN, AND J. VAN REENEN (2013): “Identifying technology spillovers and product market rivalry,” *Econometrica*, 81, 1347–1393.
- CELLINI, R. AND L. LAMBERTINI (2009): “Dynamic R&D with spillovers: Competition vs cooperation,” *Journal of Economic Dynamics and Control*, 33, 568–582.
- D’ASPREMONT, C. AND A. JACQUEMIN (1988): “Cooperative and noncooperative R&D in duopoly with spillovers,” *American Economic Review*, 78, 1133–1137.
- DAWID, H., M. Y. KEOULA, M. KOPEL, AND P. M. KORT (2015): “Product Innovation Incentives by an Incumbent Firm: A Dynamic Analysis,” *Journal of Economic Behavior & Organization*, 117, 411–438.
- DE BONDT, R. (1997): “Spillovers and innovative activities,” *International Journal of Industrial Organization*, 15, 1–28.
- FALCONE, M. (2003): “An Introduction to Semi-Lagrangian Schemes for Second Order Hamilton–Jacobi Equations,” in *RIMS Kôkyûroku*, Kyoto, Japan: Research Institute for Mathematical Sciences, Kyoto University, vol. 1323, 124–146.
- FALCONE, M. AND R. FERRETTI (2013): *Semi-Lagrangian Approximation Schemes for Linear and Hamilton-Jacobi Equations*, Society for Industrial and Applied Mathematics.
- FREIDLIN, M. I. AND A. D. WENTZELL (2012): *Random Perturbations of Dynamical Systems*, Grundlehren der mathematischen Wissenschaften, Berlin, Heidelberg: Springer, 3 ed.
- GRIFFITH, R., S. REDDING, AND J. VAN REENEN (2004): “Mapping the Two Faces of R&D: Productivity Growth in a Panel of OECD Industries,” *Review of Economics and Statistics*, 86, 883–895.
- HINLOOPEN, J., G. SMRKOLJ, AND F. WAGENER (2013): “From mind to market: A global, dynamic analysis of R&D,” *Journal of Economic Dynamics and Control*, 37, 2729–2754.
- JAFFE, A. B. (1986): “Technological Opportunity and Spillovers of R&D: Evidence from Firms’ Patents, Profits, and Market Value,” *American Economic Review*, 76, 984–1001.
- KAMIEN, M. I., E. MULLER, AND I. ZANG (1992): “Research Joint Ventures and R&D Cartels,” *American Economic Review*, 82, 1293–1306.
- KAMIEN, M. I. AND I. ZANG (2000): “Meet Me Halfway: Research Joint Ventures and Absorptive Capacity,” *International Journal of Industrial Organization*, 18, 995–1012.
- KATSOU LACOS, Y. AND D. ULPH (1998): “Endogenous spillovers and the performance of research joint ventures,” *The Journal of Industrial Economics*, 46, 333–357.

- LEAHY, D. AND J. P. NEARY (2007): “Absorptive Capacity, R&D Spillovers, and Public Policy,” *International Journal of Industrial Organization*, 25, 1089–1108.
- LIAO, C. T. (2020): “Laggards Imitate, Leaders Innovate: The Heterogeneous Productivity Effect of Imitation versus Innovation,” *Industrial and Corporate Change*, 29, 375–394.
- LIVANIS, G. AND A. LAMIN (2016): “Knowledge, Proximity and R&D Exodus,” *Research Policy*, 45, 8–26.
- MARTIN, S. (2002): “Spillovers, Appropriability, and R&D,” *Journal of Economics*, 75, 1–32.
- MOTTA, M. (2004): *Competition Policy: Theory and Practice*, Cambridge: Cambridge University Press.
- ORLANDO, M. J. (2004): “Measuring Spillovers from Industrial R&D: On the Importance of Geographic and Technological Proximity,” *The RAND Journal of Economics*, 35, 777–786.
- PETIT, M. L. AND B. TOLWINSKI (1999): “R&D Cooperation or Competition?” *European Economic Review*, 43, 185–208.
- POYAGO-THEOTOKY, J. (1999): “A Note on Endogenous Spillovers in a Non-Tournament R&D Duopoly,” *Review of Industrial Organization*, 15, 253–262.
- SALANT, S. W. AND G. SHAFFER (1998): “Optimal Asymmetric Strategies in Research Joint Ventures,” *International Journal of Industrial Organization*, 16, 195–208.
- SAMANO, M., M. SANTUGINI, AND G. ZACCOUR (2017): “Dynamics in research joint ventures and R&D collaborations,” *Journal of Economic Dynamics and Control*, 77, 70–92.
- SMRKOLJ, G. AND F. WAGENER (2019): “Research among copycats: R&D, spillovers, and feedback strategies,” *International Journal of Industrial Organization*, 65, 82–120.
- SUZUMURA, K. (1992): “Cooperative and Noncooperative R&D in an Oligopoly with Spillovers,” *American Economic Review*, 82, 1307–1320.