

Technology Ladders and R&D in Dynamic Cournot Markets

Mike Ludkovski* Ronnie Sircar†

June 2015, revised November 16, 2015

Abstract

Motivated by dramatic and unpredictable technological advances in energy production (for instance drilling and extraction of shale oil), we extend Cournot models of competition to incorporate research and development (R&D) that can lead to (stochastic) drops in production costs. Our model combines features of patent racing with dynamic market structure, capturing the interplay between the immediate competition in terms of production rates and the long-term competition in R&D. The resulting Markov Nash equilibrium is found from a sequence of one-step static games arising between R&D successes, and several numerical examples and extensive analysis of the emerging comparative statics are presented. Analyzing the relationship between current market dominance and the level of R&D investments, we find that market leaders tend to invest more, which in some sense makes oligopoly dynamically unstable. We show that anticipated market transitions have long-term impact; for example the potential of future monopoly can spur R&D investment now, even if the firm is presently uncompetitive and not actively producing. We also show that, surprisingly, random innovations have an ambiguous effect on R&D. This feature, which is driven by the Cournot framework, contrasts with the common situation whereby uncertainty lowers innovation and delays R&D investments. Finally, we demonstrate that increased competition may actually increase efforts to innovate through higher desire to achieve dominance. This would match the anecdotal evidence that the threat of market entrants forces incumbents to maintain high R&D.

Keywords: Cournot markets, R&D innovations, technology ladder, dynamic oligopoly

JEL Codes: D43, O32, C73, L13

1 Introduction

In the past decade, major advances in deep-sea offshore oil, tanker technology, and especially shale oil have fundamentally altered the oil market, marginalizing traditional producers and elevating new players. These effects can be ascribed to dramatically lower production costs of these technologies. A recent article in the New York Times quotes that “the break-even price for operating in 75 percent of the shale oil fields a year ago was \$75 a barrel, but that is now down to roughly \$60 because of innovation and lower service company costs... the break-even cost could go as low as \$50 before long.” [Krauss, 2015] The alluded innovations are uncertain and unpredictable – witness

*Department of Statistics and Applied Probability, University of California Santa Barbara, South Hall, Santa Barbara, CA 93106-3110; ludkovski@pstat.ucsb.edu. We thank Mark Duggan for research assistance, Rene Aid and Hans Tuenter for helpful discussions. We also thank participants at the Banff Workshop on New Directions in Financial Mathematics and Mathematical Economics (July 2014) and IPAM Workshop on Commodity Markets and their Financialization (May 2015) for their feedback.

†ORFE Department, Princeton University, Sherrerd Hall, Princeton, NJ, 08544; sircar@princeton.edu. Work partially supported by NSF grant DMS-1211906.

the disruptive power of shale production which was a sudden development unanticipated even 15 years ago. Conversely, consider the story of arctic oil production which remains minimal even after decades of exploration due to unforeseen challenges. Moreover, technical progress is not a one-time event, but a series of changes representing accumulation of knowledge stock and corresponding technology advances, all requiring sustained investments in research and development.

Motivated by these economic realities, in this paper we investigate dynamic stochastic Research & Development (R&D) games. The underlying framework of a non-cooperative oligopoly is a popular tool for describing commodity market equilibrium and is to be understood broadly, for example via various producer types in the oil market (conventional sweet crude, offshore oil, oil sands, shale oil, etc.) The dynamic aspect arises due to the two time scales for competition between producers. In the short term, firms compete on quantity, interacting through the aggregate supply-demand equilibrium. In the long term, firms also compete on innovation through generating structural competitive advantages. As described above for the crude oil market, one way to gain advantage is to improve efficiency through lower extraction costs.

To capture game effects which are present on both time scales we consider a Cournot market model with the producers endogenously improving their production costs. More precisely, we describe technological advances as a controlled point process, where the timing of innovation events is influenced by the investment in R&D. Consequently, progress is stochastically dependent on the research effort but is totally unpredictable otherwise. Innovation gains are assumed to be permanent and private, generating a durable competitive advantage to the innovator. The discrete nature of innovation implies that the market evolves through a sequence of technology *stages*. This setup also decouples the instantaneous equilibrium, equivalent to a static Cournot game, from the long-term cost competition that is described through dynamic-programming-type recursions.

1.1 Contributions

Our dynamic R&D game provides a link between models of stochastic innovation/patent races and of Cournot markets. Integrating both aspects in a single game theoretic framework gives insights into the interplay between joint optimization of production and R&D. Notably, the developed model endogenizes the market structure, making the number of active producers time-varying, so as to allow an endogenous transition between, say, duopoly and monopoly. Thus, we are able to capture the aforementioned history of oil production whereby some technologies dynamically enter and leave the market (e.g. the recent entry of shale production and the resulting effective suspension of deep-sea exploration). Endogenous market structure is only possible in a multi-stage setup and highlights the dynamism of our framework.

Our analysis yields several novel insights and features. First, we analyze the relationship between current market dominance and the level of R&D investments. We find that market leaders tend to invest more, which in some sense makes oligopoly dynamically unstable. Second, we show that anticipated market transitions have long-term impact; for example the potential of future monopoly can spur R&D investment now, even if the firm is presently uncompetitive and not actively producing. Third, we provide results and extensive discussion on the role of stochasticity in R&D. The fact that innovations are fundamentally uncertain in our model implies that there is always a range of *scenarios* for the future competition conditions. Surprisingly we show that random innovations have an ambiguous effect on R&D. This feature, which is driven by the Cournot framework, contrasts with the common situation whereby uncertainty lowers innovation and delays R&D investments. Fourth, we investigate the role of competition in R&D. We demonstrate that increased competition may actually increase efforts to innovate through higher desire to achieve dominance. This would match the anecdotal evidence that the threat of market entrants forces incumbents to maintain high R&D.

On the modeling front, our setup is mathematically tractable, directly building on top of the classical theory of Cournot competition. As such, we are able to use insights from *static* games to shed light on features of the dynamic model. In particular, we analyze the relationship between the static dependence of profits on costs and the dynamic R&D investments. Furthermore, our model is also highly flexible, and can be easily modified to consider further interactions between players in the R&D space, including R&D spill-overs, exogenous technical shocks or complementarity between R&D and production expenditures. This adds a new dimension to the standard Cournot competition based solely on quantity. Lastly, the multi-stage aspect of R&D innovations allows for straightforward numerical implementation, breaking down the stochastic differential game into a sequence of simple (nonlinear) optimization problems.

As explained, our main motivation comes from commodity markets where Cournot oligopoly models are well-established. Nevertheless, our model is broadly applicable since any sustained economic growth requires ongoing productivity gains. Thus, the associated technological progress and the resources allocated to it are a crucial ingredient in generic economic growth models. As detailed below, our framework can therefore be transferred to other settings, for example dynamic races between consumer technology firms (the Apple vs. Samsung archetype) or multi-stage patent races.

1.2 Related Frameworks

The main ingredients of our model are (i) Cournot competition, (ii) endogenous R&D effort with stochastic innovation, and (iii) continuous-time dynamic framework. The precise combination of these building blocks appears here for the first time, but there are several substantial bodies of relevant research.

First, there is a large strand of literature studying the monopolist’s (or social planner’s) problem of profit maximization under endogenous technological progress. This is usually done with a Ramsey-type economic growth model that combines the production and R&D activities within a single framework. Historically, this analysis originated in the study of exhaustible resource extraction [Kamien and Schwartz, 1978, Pindyck, 1980]. In endogenous growth models, R&D is used to increase the knowledge stock, which in turn raises production efficiency. See for example Goulder and Schneider [1999], Lafforgue [2008], Grimaud et al. [2011]. One recent key topic involves R&D to develop a renewable energy backstop to guard against exhaustibility of conventional fossil fuels [Tsur and Zemel, 2003]. Our model can be seen as a game-theoretic expansion of these ideas. Rivalry considerations add a new dimension to investing in R&D, because the value of R&D is inextricably tied to the size of competition which in turn is driven by innovations.

In terms of the R&D innovations model, our setup is closest to the work of Lafforgue [2008]. Lafforgue considered a central planner framework that features a single consumption good and a representative consumer. Producing the consumption good requires labor and has an efficiency parameter $B(t)$. Lafforgue assumed that $B(t)$ can be sequentially improved through R&D expenditures. As below, Lafforgue [2008] represents R&D innovations as a controlled counting process where each innovation increases $B(t)$ by a factor of $1 + b$. In addition, Lafforgue [2008] assumes an exhaustible resource with stock $X(t)$ that is required for production, and a perfect substitutability between production and R&D labor. By postulating specific analytic forms of the corresponding dynamics, production and utility functions, an explicit analytic solution is then obtained to the stochastic optimization problem. Our model can be seen as the generalization of this setup to the multiple-agent game setting. With multiple agents, the central planner perspective no longer applies, and new dynamic features (most notably different game regimes in terms of player participation) emerge. The increased complexity does rule out analytic solutions.

From the control perspectives, all the above models can be first classified as deterministic or

stochastic, and second as using continuous, singular or impulse (stopping) controls. In the context of R&D development, because innovations are typically indivisible or lumpy, the common stochastic source is Poissonian (in contrast, for capacity expansion or for modeling demand fluctuation, Brownian shocks are usually used). In our model below, we also use Poissonian shocks but work with continuous controls, which allows us to leverage the tractable framework of piecewise-deterministic control [Davis, 1993].

Moving beyond central-planner models, R&D has been viewed in the context of patent races, see e.g. the survey by Reinganum [1989]. In the classical patent race players compete to receive a single prize by making R&D investments that stochastically determine the winner. Starting from the basic one-shot framework of Reinganum [1983], multi-stage extensions have been analyzed in Grossman and Shapiro [1987], Harris and Vickers [1987], Judd [2003] and Doraszelski [2003]. With multiple stages, asymmetry between agents become central to the analysis. The typical outcome is of “increasing dominance” —the current leader that is closest to the prize also puts in more effort, which is essentially driven by the increased probability of collecting the prize, the so-called “pure progress effect”. The above models assume only a single reward that is fully appropriated by the winner; in contrast we embed a patent race within a competitive market framework. The latter makes the R&D race less of a zero-sum and moreover introduces further effects due to dynamic market structure.

Alternatively, R&D race can be viewed as a timing game, where sustained R&D efforts are replaced with a one-shot investment. Starting with the seminal work of Fudenberg and Tirole [1985], there has been ongoing interest in such preemption games that generalize the real options setting to multiple agents. Stochastic game models under a variety of uncertainty models (fluctuating demand, exogenous or endogenous technology shocks, duopoly or oligopoly, etc.) have been considered, see Weeds [2002], Huisman and Kort [2004], Femminis and Martini [2011] or the recent review in Azevedo and Paxson [2014]. A major topic in this research is to determine the market structure. However, because the games are one-shot (the only action is the timing of investment) this market structure is essentially static *a priori*. In contrast, below we work with a genuinely dynamic setting where agent strategies are ongoing and adaptive and where the different game stages allow for percolation of different equilibrium types over time. Stopping-time games are both simpler (since strategies can be typically summarized in terms of a simple threshold-exercise rule) and more complex (allowing for e.g. both sequential and simultaneous exercise) than a dynamic game.

Better technology confers first-mover advantage and hence is naturally linked to leader-follower games. A notable reference is Fölster and Trofimov [1997] who considered an oligopoly where each of n firms maximizes R&D effort $a(t)$ that stochastically determines the random first innovator who temporarily collects extra profits. While Folster and Trofimov have a similar “quality ladder” for cost reduction within a Cournot market, their model is effectively one-stage in our notation and hence has (an endogenous) static market structure.

Lastly, while the model herein is intrinsically stochastic, it bears resemblance to deterministic models of R&D races, such as the discrete-stage frameworks of Fudenberg et al. [1983] or the continuous-time differential game of Cellini and Lambertini [2009]. The main difference is that a deterministic model leaves little place for asymmetry since equilibrium is fully determined by the initial condition. Thus, typically only the completely symmetric case is interesting, which is known as ε -preemption [Fudenberg et al., 1983]. In turn, symmetry destroys much of the game aspect: for example in Cellini and Lambertini [2009] the completely symmetric equilibrium is analytically identical to classical optimization. In our model the uncertainty plays a fundamental role since it allows for varied dynamic market structures.

To sum up, our work is at the intersection of three research streams. First, we extend the ongoing stochastic innovations model of Lafforgue [2008] to the game setting. Rivalry consideration are

crucial since they modify market structure and induce additional strategic considerations. Second, we combine the multi-stage patent races [Grossman and Shapiro, 1987, Judd, 2003] with a Cournot market. Third, we extend the one-shot stochastic games of Fölster and Trofimov [1997] to the dynamic setting. New dynamic effects include anticipation of blockading and possibility to analyze the time-profile of R&D across no-longer symmetric players.

The rest of the paper is organized as follows. Section 2 develops the dynamic Cournot market model we employ, in particular describing technology ladders capturing cost improvements. Section 3 constructs a Markov Nash equilibrium for the above model by leveraging the local structure of static Cournot market and the Poissonian innovation process. The resulting endogenized technical progress and respective market structure are investigated in Section 4 for the unilateral R&D case, and in Section 5 for the bilateral R&D. Section 7 addresses several extensions of the basic model, including partial substitutability, spillovers and deterministic R&D. Section 8 concludes, while the Appendix contains proofs of several key lemmas.

2 Cournot Oligopoly with Technology Innovation

We consider a Cournot oligopoly with $L \geq 2$ players or producers. (The monopoly case $L = 1$ also obeys the properties below, but in this paper we are primarily concerned with markets where there is competition.) The players compete in a single market by choosing their production rates q_i . Equilibrium emerges based on a demand curve $D(\cdot)$ and aggregate supply

$$Q = q_1 + \dots + q_L.$$

Thus, the market clearing price P received by each producer is determined according to

$$P = D^{-1}(Q).$$

Note that the above assumes perfectly substitutable goods from different producers. This choice is to simplify the presentation and is not essential to the analysis that follows; see Section 7.2 for discussion of differentiated markets.

To maintain finite market capacity and avoid other technical difficulties, the next assumption imposes regularity on the relationship between P and Q .

Assumption 1. $Q \mapsto P(Q)$ is twice continuously differentiable with $P'(Q) < 0$ everywhere, and there exists $\eta < \infty$, such that $P(\eta) = 0$. Moreover, $QP(Q)$ is bounded from above.

The upper limit η is called the saturation point – total production Q is guaranteed to stay below η , otherwise prices will collapse to zero. The next further assumption restricts the convexity of the price function $P(Q)$ and will be used in the sequel to explicitly describe the Cournot oligopoly equilibrium. Define

$$\rho(Q) := \frac{-QP''(Q)}{P'(Q)}$$

to be the relative prudence of the price function.

Assumption 2. The price function $P(Q)$ satisfies $\sup_{Q \in (0, \eta)} \rho(Q) =: \bar{\rho} < 2$.

Notation: we use the generic subscript i, j to denote a particular player $i = 1, \dots, L$; the multi-index $-i = \{1, 2, \dots, i-1, i+1, \dots, L\}$ denotes all players except i . In particular, the aggregate supply consists of player i 's production and the rest, $Q = q_i + Q_{-i}$. When $L = 2$ we use $j \neq i$ to denote the other player, so that $Q = q_i + q_j$.

Players are differentiated according to their production costs $C_i \geq 0$, which linearly enter into their profit rate

$$\pi_i = q_i \cdot (P(Q) - C_i). \quad (1)$$

In the short-term, players optimize π_i by establishing a Nash equilibrium in terms of production rates q_i ; interactions are solely through the clearing price $P(Q)$. In the long-term players have the ability to lower their production costs through Research and Development activities. Thus, $C_i = C_i(t)$ may change over time; lower costs will translate into higher profits π_i in (1).

We assume that technology changes are abrupt and possible stages are summarized by a *technology ladder*

$$\mathcal{C}_i := \{c_i(n), n = 1, \dots, \}, \quad c_i(n) \geq c_i(n+1) \geq \dots \geq 0. \quad (2)$$

At any given date t , $C_i(t) \in \mathcal{C}_i$ is discrete, and can be encoded via the corresponding stage n . The ladder \mathcal{C}_i may contain either finite or infinite number of stages and is fixed a priori. Note that costs are required to stay non-negative, so C_i ranges in the finite interval $[0, c_i(1)]$. Also the ladder is fixed for all time in our models to enable us to construct a time stationary solution. In general, it is possible to treat randomly evolving or time-dependent ladders if one is willing to incorporate more state variables into the game functions.

Two illustrative examples are

- Linear progress $c_i(n) = 1 - \mu n$, $n = 1, 2, \dots, \lfloor 1/\mu \rfloor$;
- Exponential progress $c_i(n) = \exp(-\mu n)$, $n \in \mathbb{N}$.

The first case corresponds to a fixed absolute improvement in costs and has a finite number of stages to keep c_i positive; the second case to a percentage improvement of $\mu\%$ with each new technology advance. Both examples start with $c_i(0) = 1$ and approach 0 as n grows. As we show below, under mild assumptions this also makes \mathcal{C}_i effectively finite, as far as the game is concerned.

2.1 Technology Innovation Process

The role of R&D is to induce progress by moving up to higher stages of technology along the ladder in (2). For simplicity, we assume that each technical innovation moves the corresponding producer i one step up the ladder, lowering her costs from $c_i(n)$ to $c_i(n+1)$. See Judd [2003] for more general descriptions. Let $N_i(t) \in \mathbb{N}$ be the index of technological progress of player i at epoch t , so that $C_i(t) = c_i(N_i(t))$. The overall state of progress is then summarized by the state vector $\mathbf{N}(t) \equiv (N_1(t), \dots, N_L(t)) \in \mathbb{N}^L$. Technical progress is uncertain, i.e. $(\mathbf{N}(t))$ is a stochastic process.

As explained, innovations are “lumpy” or discrete, i.e. $t \mapsto N_i(t)$ is piecewise constant in time, and can be thought of as a counting process. To endogenize R&D investments we link the *hazard* rate $\lambda_i(t)$ of N_i with the R&D effort levels $a_i(t)$ that are continuously controlled by the players. Denote by $\mathbf{a}(t)$ the profile of R&D efforts. We postulate

$$\lambda_i(t) = \bar{\lambda} a_i(t) \quad (3)$$

so that $\lambda_i(\cdot)$ is linear in effort. The scaling constant $\bar{\lambda}$ modulates the overall speed of innovation, see Section 6. Since the measurement units of effort are arbitrary, the linear link in (3) is effectively without loss of generality. However, note that (3) implies that innovation for player i is independent of either the behavior of other players, or their present technological state $\mathbf{N}(t)$. In Section 7.4 we discuss a more general version of (3) that allows for various spillover effects.

R&D expenditures are costly: effort at level a_i carries a running cost $\mathcal{R}_i(a_i)$ to player i . In line with the concept of diminishing returns, we assume

Assumption 3. The R&D cost function $a \mapsto \mathcal{R}_i(a)$ is differentiable, convex, with $\mathcal{R}_i(0) = 0$ and $\lim_{a \rightarrow \infty} \mathcal{R}'_i(a) = +\infty$.

Convex costs guarantee that efforts remain finite which in turn ensures that N_i is well-defined (i.e. does not explode). We define $\tau_i^n := \inf\{t : N_i(t) \geq n\}$ to be the time that player i achieves innovation stage n . The above assumptions imply that $\mathbb{P}(\tau_i^n < \tau_i^{n+1} < \infty, \forall n) = 1$, so that instantaneous innovations are ruled out, and \mathbf{N} is a bona fide multi-dimensional counting process.

Remark 1. An alternative rephrasing of our model harking back to Kamien and Schwartz [1978] is based on the idea that innovation is brought about through cumulating knowledge stock. Given \mathbf{a} and starting with (3), define $\Lambda_i(t) := \int_0^t \lambda_i(s) ds$ to be the cumulative knowledge investment. Let I_1, I_2, \dots , be a sequence of random variables with unit exponential increments, $I_{k+1} - I_k \sim \text{Exp}(1)$, that are independent of Λ_i . Define $\tilde{N}_i(t) = \sup\{k : I_k \leq \Lambda_i(t)\}$. Then \tilde{N}_i is equivalent to our model of the innovation process N_i . Thus, growing knowledge randomly triggers the occurrence of an innovation (note that I_k 's have the same distribution as the arrival times of a standard Poisson process) and the conditional probability of successful technological change $\mathbb{P}(N_i(t) > k | \Lambda_i(t))$ is a strictly increasing and known function of $\Lambda_i(t)$. There is also extensive literature [Reinganum, 1989, Doraszelski, 2003] on knowledge accumulation races, in particular allowing for further features such as learning-by-doing, knowledge decay, etc. Here we stick to the ‘‘Markov’’ knowledge frameworks which summarize progress in terms of the technology state \mathbf{N} .

A set of strategies is thus described by the $2L$ -dimensional process (\mathbf{q}, \mathbf{a}) which specifies for each player her continuous production rate $q_i(t) \geq 0$ and her continuous effort level $a_i(t) \geq 0$. Overall, \mathbf{q} drives the instantaneous profit streams received by the players, while \mathbf{a} modulates the technological innovations summarized by \mathbf{N} . Similar to \mathbf{N} , (\mathbf{q}, \mathbf{a}) will be dynamic and can be chosen by players to adapt to the randomly evolving market state $\mathbf{N}(t)$.

2.2 R&D Game Formulation

Since $N_i(t)$ is stochastic, future production costs and hence profits are uncertain. Players evaluate their expected total discounted future profits via the performance measure

$$J_i(\mathbf{n}; \mathbf{q}, \mathbf{a}) := \mathbb{E} \left[\int_0^\infty e^{-\rho_i t} \left[q_i(t) (P(Q(t)) - c_i(N_i(t))) - \mathcal{R}_i(a_i(t)) \right] dt \mid \mathbf{N}(0) = \mathbf{n} \right], \quad (4)$$

where $\rho_i > 0$ is the intertemporal discount factor of player i . The expectation in (4) is with respect to the random shocks embedded in $(\mathbf{N}(t))$, and J_i is a function of the strategy profile (\mathbf{q}, \mathbf{a}) , and the initial conditions $N_i(0) = n_i, i = 1, \dots, L$.

The above model (4) yields a dynamic stochastic noncooperative game with players aiming to maximize J_i and interacting through the joint price received $P(Q(t))$. To describe the resulting equilibrium we rely on the notion of Markov perfect Nash equilibrium (see, for instance, [Vives, 2001]), so that all strategies are functions of the main state $\mathbf{N}(t)$. Hence, given $\mathbf{N}(t) = \mathbf{n}$, each player looks for an action $(q_i(t), a_i(t))$ which maximizes her net present value J_i treating the other players as fixed. Note that q_i is only used for the immediate Cournot market equilibrium, while a_i controls the intertemporal transitions in \mathbf{N} .

Since lower costs are associated with higher profits (see Section 2.4), players have an incentive to invest in R&D. However, the stochastic nature of R&D and Cournot effects make these investments ambiguous. First, innovation success is not guaranteed, so players must average over the potential future scenarios. Second, like in classical patent races, simultaneous R&D investments by the players produce random ‘‘winners’’, so that the role of the technology leader is not fixed. Third, changing production costs $C_i(t)$ affect the market structure and in particular the number of players

actively producing. As a result, R&D introduces *phase transitions* (e.g. from oligopoly to monopoly or vice-versa) making the respective impacts on profits nonsmooth.

Remark 2. Above we take R&D effort as sustained, i.e. continuous in time. A complementary view is linked to the concept of capacity expansion and views $a(t)$ as instantaneous control. In that case, $a_i(t)$ is series of actions and can be modeled via the framework of multiple optimal stopping, e.g. Dahlgren and Leung [2015]. See Goyal and Netessine [2007] for models of oligopolistic capacity expansion (e.g. between competing real estate developers).

2.3 Mathematical Details

To formalize the concept of equilibrium under a strategy profile we briefly recall the construction of controlled point processes necessary to describe the component-wise dynamics of \mathbf{N} [Bremaud, 1981]. Let Ω be the space of L -dimensional counting process paths, i.e. all piecewise-constant, increasing, right-continuous paths $\omega = (\omega(t))$, $t \in \mathbb{R}_+$, $\omega(t) \in \mathbb{N}^L$ satisfying $\omega_i(t) - \omega_i(t-) = 0$ or 1 for all t . Define the canonical processes $N_i(t; \omega) = \omega_i(t)$, and let (\mathcal{F}_t) be the natural filtration of the aggregated $(\mathbf{N}(t))$, $\mathcal{F}_t := \sigma(\mathbf{N}(s) : s \leq t)$, with $\mathcal{F} = \mathcal{F}_\infty$. Every measure P on $(\Omega, \mathcal{F}, (\mathcal{F}_t))$ is described through its compensator process $(\Lambda(t))$. We restrict attention to absolutely continuous and bounded compensators which implies the existence of a hazard rate process λ_i satisfying $\Lambda_i(t) = \int_0^t \lambda_i(s) ds$ for each $i = 1, \dots, L$. The above assumption means that $t \mapsto N_i(t) - \int_0^t \lambda_i(s) ds$ is a (\mathcal{F}_t) -martingale and the behavior of N_i is completely specified through $\lambda_i(t)$. Moreover, it implies that given $\lambda_1(t), \dots, \lambda_L(t)$, the different components N_i 's are conditionally independent. The conditional independence of N_i 's means that there is no interaction between innovation processes beyond the hazard rates.

The above construction allows to assign rigorous meaning to (3) for any (\mathcal{F}_t) -adapted R&D strategy profile \mathbf{a} . We similarly consider (\mathcal{F}_t) -adapted production strategies \mathbf{q} . Finally, after specifying the initial condition $\mathbf{N}_0 = \mathbf{n}$, this allows to assign a rigorous meaning to the probability measure $\mathbb{P}_{\mathbf{n}}$ that appears in (4) (for example this can be done using a change-of-measure technique starting from a reference measure \mathbb{P}^0 , see Bremaud [1981]).

2.4 Effect of Production Costs in Static Cournot Games

To explain the effects of R&D under dynamic equilibria, we briefly recall the role of production costs in *static* Cournot games. Consider a static Cournot game with one-shot payoff

$$\pi_i = q_i(P(q_i + Q_{-i}) - c_i), \quad \text{for player } i = 1, 2, \dots, L.$$

Each player chooses a nonnegative production level q_i , competing solely through the aggregate production $Q = \sum_i q_i$. The game will be indexed by the cost profile $\mathbf{c} = (c_1, \dots, c_L)$ where, without loss of generality, we order the players by increasing production costs $c_1 \leq c_2 \leq \dots$, and define for any $\ell \leq L$,

$$B_\ell = c_1 + \dots + c_\ell. \tag{5}$$

We first recall the corresponding equilibrium theory from Harris et al. [2010]. See also the textbook treatment of Cournot games in Vives [2001]. Under a Nash equilibrium $q^*(\mathbf{c}) \equiv (q_1^*, \dots, q_L^*)$ we expect the candidate production rates to satisfy the first order conditions

$$\frac{\partial \pi_i}{\partial q_i} \Big|_{q_i=q_i^*} = 0 \quad \Leftrightarrow \quad q_i^* \cdot \frac{\partial P(q_i^* + Q_{-i}^*)}{\partial q_i} + P(Q^*) = c_i. \tag{6}$$

However, the constraint $q_i^* \geq 0$ may be binding, so that the number of players actively participating in the equilibrium (i.e. those with $q_i^* > 0$ strictly positive) is to be determined. Moreover, existence and uniqueness of the above equilibrium must be determined. Players that have $q_i^* = 0$ are said to

be blockaded—they are not producing and collect zero profits. Because players are sorted in terms of their production costs, the active set $\{i : q_i^* > 0\}$ is of the form $i = 1, 2, \dots, \ell^*$ for some $\ell^* \leq L$.

Under Assumption 2, the next Proposition from Harris et al. [2010] explicitly characterizes the resulting unique Nash equilibrium.

Proposition 1. [Harris et al., 2010, Lemma 2.3]. *For $\ell = 1, 2, \dots, L$, let*

$$f_\ell(Q) = QP'(Q) + \ell P(Q),$$

and let Q_ℓ be the unique solution in $(0, \eta)$ of $f_\ell(Q_\ell) = B_\ell$, where we defined B_ℓ in (5). Take $Q^ = \max\{Q_\ell : 1 \leq \ell \leq L\}$. Then the unique Nash equilibrium production rates are*

$$q_i^*(\mathbf{c}) = \max\left\{\frac{P(Q^*) - c_i}{-P'(Q^*)}, 0\right\}, \quad (7)$$

and the number of players producing in equilibrium is $\ell^ = \min\{\ell : Q^* = Q_\ell\}$. Moreover the aggregate production is exactly Q^* :*

$$Q^* = Q^*(\mathbf{c}) = \sum_{i=1}^L q_i^*(\mathbf{c}).$$

Proposition 1 completely characterizes the production controls for players with asymmetric but fixed costs. Note that it relies strongly on Assumption 2 which is a sufficient condition to guarantee existence of equilibrium; if $\bar{\rho} > 2$, there might be no equilibrium.

In the model below we will let \mathbf{c} change and so it is important to understand the comparative statics on

$$\pi_i^*(\mathbf{c}) = q_i^*(\mathbf{c}) (Q^*(\mathbf{c}) - c_i). \quad (8)$$

We assume that we are in the “normal” case whence a decrease in player- i costs c_i raises her production and profit, and lowers the production and profits of the other players. The general link is based on the stability of the equilibrium which in turn is tied to the shape of the best-response curves $q_i^*(Q)$ in (6). We refer to [Vives, 2001, Section 4.3] for a discussion of this problem, including examples where higher costs *increase* production for all players. As a summary we have the following

Proposition 2. *Suppose that $P'(Q^*) + q_i^* P''(Q^*) \leq 0 \forall i$ in the equilibrium \mathbf{q}^* . Then*

$$\frac{\partial \pi_i^*}{\partial c_i} < 0 \quad \text{and} \quad \frac{\partial \pi_i^*}{\partial c_j} > 0, \quad \text{for all } i, j.$$

Constant Prudence Price Functions

A typical parametric family of price functions is given by the constant-prudence $\rho(Q) \equiv \rho$ inverse demand curves

$$P(Q) = \begin{cases} \frac{\eta}{1-\rho} \left(1 - \left(\frac{Q}{\eta}\right)^{1-\rho}\right) & \rho \neq 1; \\ \eta(\log \eta - \log Q) & \text{if } \rho = 1. \end{cases} \quad (9)$$

As before, η is the choke price. For such price functions we can explicitly describe the unique Nash equilibrium and the impact of the production costs on profits. In this case Assumption 2 is not necessary for existence of equilibrium and we can take ρ such that $\rho < L + 1$.

Proposition 3. [Harris et al., 2010, Proposition 2.9]. *Assume the price function $P(Q)$ is in the constant prudence family (9), with $\rho < L + 1$. Then there is a unique Nash equilibrium given by*

$$q_i^*(\mathbf{c}) = \left(\frac{Q^*}{\eta}\right)^\rho \max(\bar{P} - c_i, 0) \quad (10)$$

with

$$\bar{P} = \min_{1 \leq \ell \leq L} \frac{B_\ell + \eta}{\ell + 1 - \rho}, \quad \text{and} \quad Q^* = \begin{cases} \eta \left(1 - (1 - \rho) \frac{\bar{P}}{\eta}\right)^{\frac{1}{1-\rho}} & \rho \neq 1; \\ \eta \exp\left(-\frac{\bar{P}}{\eta}\right) & \rho = 1, \end{cases} \quad (11)$$

and we recall from (5) that B_ℓ is the sum of the costs of the first ℓ players. The number of active players is $\ell^* = \max\{\ell : \eta + B_{\ell-1} - (\ell - \rho)c_\ell > 0\}$. The corresponding profit is

$$\pi_i^*(\mathbf{c}) = q_i^*(\mathbf{c})(\bar{P} - c_i).$$

The next lemma also shows that the constant-prudence family enjoys the intuitive comparative statics in \mathbf{c} . Its proof is given in the Appendix.

Lemma 1. *Suppose the price function $P(Q)$ has a constant relative prudence $\rho(Q) \equiv \rho < 2$. Then when $\ell^* \leq L$ players $i = 1, 2, \dots, \ell^*$ are active in equilibrium we have*

$$\frac{\partial q_i^*}{\partial c_j} = \begin{cases} \frac{1}{P'(Q^*)} \left[\frac{\ell^* - (1 - \frac{q_i^*}{Q^*})\rho}{\ell^* + 1 - \rho} \right] & \text{if } j = i < \ell^*; \\ \frac{1}{P'(Q^*)} \left[\frac{-1 + \frac{q_i^*}{Q^*}\rho}{\ell^* + 1 - \rho} \right] & \text{if } j \neq i < \ell^*; \\ 0 & \text{otherwise.} \end{cases} \quad (12)$$

Therefore, for $i, j < \ell^*$, we have that $\frac{\partial q_i^*}{\partial c_i} < 0$ and $\frac{\partial q_i^*}{\partial c_j} > 0$.

Lemma 1 connects cost sensitivity to the fraction of total production attributable to player i , $\frac{q_i^*}{Q^*}$. We recall from Harris et al. [2010] that this quantity is given by :

$$\frac{q_i^*}{Q^*} = \frac{P(Q^*) - c_i}{\sum_{j=1}^L (P(Q^*) - c_j)}.$$

Therefore, fixing Q^* , players that have a higher production rate q_i^* or a lower cost c_i are more sensitive to both their own costs and competitor costs. Thus, current market leaders (in terms of market share) are also the ones that have the strongest (static) incentive to engage in R&D.

In the rest of the paper, we analyze the dynamic game introduced in Section 2.2.

3 Dynamic Equilibrium

As our notion of equilibrium we focus on Markov subgame perfect equilibrium (MPE) [Vives, 2001] which employs the concept of Nash equilibrium in a dynamic setting. In short, MPE assumes that equilibrium strategies are in closed-loop feedback form, so that $q_i(t) = q_i(\mathbf{N}(t))$, $a_i(t) = a_i(\mathbf{N}(t))$. Thus, players base their decisions on the aggregate market state, without any further randomization

or deviation. As explained in Section 2.3, \mathbf{N} is a counting process, i.e. $t \mapsto \mathbf{N}(t)$ is piecewise constant. It follows that so are the feedback equilibrium controls. This allows to decompose the global game into a sequence of static games indexed by \mathbf{n} .

Beyond the feedback structure, the set of admissible strategies is locally specified through the restriction $(q_i, a_i)(\mathbf{n}) \in \mathcal{A}_i(\mathbf{n})$. In our main setup we assume that production and R&D are fully independent so that any non-negative controls are admissible, $\mathcal{A}_i(\mathbf{n}) = \mathbb{R}_+^2 \forall \mathbf{n}$. This corresponds to the economic reality of R&D effort financed by *capital* [Kamien and Schwartz, 1978]. In other words, when choosing R&D levels, the only consideration is the total expected innovation gain vis-a-vis the expected revenues; the corresponding expenditures are financed by borrowing against future earnings as needed. In the context of a competitive Cournot market, capital-financed R&D seems reasonable since companies are generally able to adaptively expand or shrink their business lines independently. An alternative assumption of *labor-consuming* R&D [Lafforgue, 2008] that generates complementarity between q_i and a_i is discussed in Section 7.3. Coupling production and R&D expenditures makes \mathcal{A}_i a strict subset of the positive quadrant.

Below we explicitly construct a MPE by developing first-order-condition equations for $(\bar{q}, a^*)(\mathbf{n})$. (We have used the notation \bar{q} for the equilibrium production in the dynamic game to distinguish from the Nash equilibrium production functions q_i^* in the static game of Section 2.4). For typographical convenience we focus on the duopoly setup $L = 2$, writing out explicitly $\mathbf{n} = (n_1, n_2)$ and labeling the players as P1, P2. We also write $P(q_i, q_j) = P(q_i + q_j) = P(Q)$. The extension to oligopoly is straightforward and requires only typographical substitution, and similarly for the monopoly case $L = 1$.

Suppose that $\bar{q}_i(\mathbf{n}), a_i^*(\mathbf{n}), i = 1, 2, \mathbf{n} \in \mathbb{N}^2$ are an MPE strategy profile. We denote the corresponding game functions by

$$v_i(n_1, n_2) = \mathbb{E}_{n_1, n_2} \left[\int_0^\infty e^{-\rho t} \left[\bar{q}_i(\mathbf{N}(t)) (P(\bar{q}_i(\mathbf{N}(t)), \bar{q}_j(\mathbf{N}(t))) - c_i(N_i(t))) - \mathcal{R}_i(a_i^*(\mathbf{N}(t))) \right] dt \right]. \quad (13)$$

Next, we define the global set of admissible controls

$$\mathcal{A}_i := \{(q_i, a_i) : (q_i, a_i)(\mathbf{n}) \in \mathcal{A}_i(\mathbf{n}) \forall \mathbf{n} \in \mathbb{N}^2\}.$$

The equilibrium condition implies that

$$v_i(n_1, n_2) = \sup_{(q_i, a_i) \in \mathcal{A}_i} \mathbb{E}_{n_1, n_2} \left[\int_0^\infty e^{-\rho t} \left[q_i(\mathbf{N}(t)) \left(P(q_i(\mathbf{N}(t)), \bar{q}_j(\mathbf{N}(t))) - c_i(N_i(t)) \right) - \mathcal{R}_i(a_i(t)) \right] dt \right]. \quad (14)$$

Even though the other player's costs $C_j(t)$ do not appear in the above equation, we stress that \bar{q}_j depends on $N_j^*(t)$, so anticipating the innovations of player j is crucial for selecting the best strategy for player i .

Fix $\mathbf{N}(0) = (n_1, n_2)$ and let σ^1, σ^2 denote the first technology advance times of P1 and P2 respectively (if $a_i \equiv 0$ then $\sigma^i = +\infty$ with probability 1). Then on the random interval $[0, \underline{\sigma}]$, where $\underline{\sigma} := \sigma^1 \wedge \sigma^2$ both N_1 and N_2 are constant. The Markov property and subgame perfection

imply that

$$\begin{aligned}
v_i(n_1, n_2) &= \mathbb{E} \left[\int_0^{\underline{\sigma}} e^{-\rho_i t} [\bar{q}_i(\mathbf{N}(t))(P(\bar{q}_i(\mathbf{N}(t)), \bar{q}_j(\mathbf{N}(t))) - c_i(n_i)) - \mathcal{R}_i(a_i^*(\mathbf{N}(t)))] dt \right. \\
&\quad \left. + e^{-\rho_i \underline{\sigma}} \int_{\underline{\sigma}}^{\infty} e^{-\rho_i(t-\underline{\sigma})} [\bar{q}_i(\mathbf{N}(t))(P(\bar{q}_i(\mathbf{N}(t)), \bar{q}_j(\mathbf{N}(t))) - c_i(N_i(t))) - \mathcal{R}_i(a_i^*(t))] dt \right] \\
&= \mathbb{E} \left[\int_0^{\underline{\sigma}} e^{-\rho_i t} [\bar{q}_i(\mathbf{n})(P(\bar{q}_i(\mathbf{n}), \bar{q}_j(\mathbf{n})) - c_i(n_i)) - \mathcal{R}_i(a_i^*(\mathbf{n}))] dt \right. \\
&\quad \left. + e^{-\rho_i \underline{\sigma}} \mathbf{1}_{\{\underline{\sigma}=\sigma^1\}} v_i(n_1 + 1, n_2) + e^{-\rho_i \underline{\sigma}} \mathbf{1}_{\{\underline{\sigma}=\sigma^2\}} v_i(n_1, n_2 + 1) \right], \tag{15}
\end{aligned}$$

where the last equality follows because $N_i(s) = n_i$ is constant for $s \leq \underline{\sigma}$ for all i , so that the feedback control $a_i^*(s) = a_i^*$ is also constant. Therefore, the hazard rate of N_i is constant on $[0, \underline{\sigma}]$ and the equilibrium condition reduces to

$$\begin{aligned}
v_i(n_1, n_2) &= \sup_{(q_i, a_i) \in \mathcal{A}_i(n_1, n_2)} \mathbb{E} \left[\int_0^{\underline{\sigma}} e^{-\rho_i t} [q_i(P(q_i, \bar{q}_j(\mathbf{n})) - c_i(n_i)) - \mathcal{R}_i(a_i)] dt \right. \\
&\quad \left. + e^{-\rho_i \underline{\sigma}} \mathbf{1}_{\{\underline{\sigma}=\sigma^1\}} v_i(n_1 + 1, n_2) + e^{-\rho_i \underline{\sigma}} \mathbf{1}_{\{\underline{\sigma}=\sigma^2\}} v_i(n_1, n_2 + 1) \right], \tag{16}
\end{aligned}$$

where the supremum is now over non-negative constants q_i, a_i . Furthermore, (taking $\lambda_i = \bar{\lambda} a_i$) constant hazard rates imply that $\mathbb{P}(\underline{\sigma} > s) = \mathbb{P}(\sigma^1 > s, \sigma^2 > s) = e^{-\bar{\lambda}(a_1+a_2)s}$ so $\underline{\sigma}$ has an exponential distribution and

$$\begin{aligned}
\mathbb{E}[e^{-\rho_i \underline{\sigma}} \mathbf{1}_{\{\underline{\sigma}=\sigma^1\}}] &= \mathbb{E}[e^{-\rho_i \sigma^1} | \sigma^1 < \sigma^2] \mathbb{P}(\sigma^1 < \sigma^2) \\
&= \frac{a_1}{a_1 + a_2 + \rho_i / \bar{\lambda}}.
\end{aligned}$$

Similarly,

$$\mathbb{E} \left[\int_0^{\underline{\sigma}} e^{-\rho_i s} ds \right] = \mathbb{E} \left[\frac{1 - e^{-\rho_i \underline{\sigma}}}{\rho_i} \right] = \frac{1}{\bar{\lambda} a_1 + \bar{\lambda} a_2 + \rho_i}.$$

Plugging the above into (16), and similarly writing out the optimization problem for $v_2(n_1, n_2)$ we end this Section with the following equilibrium construction.

Proposition 4. *Suppose that for each $(n_1, n_2) \in \mathbb{N}^2$ we have value functions $v_i(n_1, n_2), i = 1, 2$ that satisfy the two-dimensional optimization system*

$$\left\{ \begin{aligned}
v_1(n_1, n_2) &= \sup_{(q_1, a_1) \in \mathcal{A}_1(n_1, n_2)} \frac{1}{\bar{\lambda} a_1 + \bar{\lambda} a_2^* + \rho_1} \left\{ q_1(P(q_1, \bar{q}_2) - c_1(n_1)) - \mathcal{R}_1(a_1) \right. \\
&\quad \left. + \bar{\lambda} a_1 v_1(n_1 + 1, n_2) + \bar{\lambda} a_2^* v_1(n_1, n_2 + 1) \right\}; \\
v_2(n_1, n_2) &= \sup_{(q_2, a_2) \in \mathcal{A}_2(n_1, n_2)} \frac{1}{\bar{\lambda} a_1^* + \bar{\lambda} a_2 + \rho_2} \left\{ q_2(P(\bar{q}_1, q_2) - c_2(n_2)) - \mathcal{R}_2(a_2) \right. \\
&\quad \left. + \bar{\lambda} a_1^* v_2(n_1 + 1, n_2) + \bar{\lambda} a_2 v_2(n_1, n_2 + 1) \right\},
\end{aligned} \right. \tag{17}$$

with the maximizers $\bar{q}_i \equiv \bar{q}_i(n_1, n_2)$, $a_i^* \equiv a_i^*(n_1, n_2)$. Then $(\bar{\mathbf{q}}, \mathbf{a}^*)$ form a MPE strategy profile with corresponding game values $v_i(\mathbf{n})$.

Remark 3. One can straightforwardly generalize the model in terms of the underlying point processes N_i . In particular, we may remove the assumption that N_i is increasing and/or has jumps of size 1 only. Economically, this corresponds to the possibility that multiple technology stages can be traversed at once, or that technology gains can be lost over time (in the sense of increasing production costs). The latter situation would be realistic for energy production where cheap conventional sources may become exhausted over time (think coal) leading to upward shocks in $C_i(t)$. Assuming that the jump distribution of N_i is independent of (q_i, a_i) and identical across $i = 1, 2$, this generalization consists of replacing (17) with

$$v_1(n_1, n_2) = \sup_{a_1, q_1} \frac{1}{\bar{\lambda}a_1 + \bar{\lambda}a_2^* + \rho_1} \left\{ \{q_1(P(q_1, \bar{q}_2, \cdot) - c_1(n_1))\} - \mathcal{R}_1(a_1) \right. \\ \left. + \bar{\lambda}a_1 \sum_k p_k v_1(n_1 + k, n_2) + \bar{\lambda}a_2^* \sum_k p_k v_1(n_1, n_2 + k) \right\}, \quad (18)$$

where $p_k := \mathbb{P}(\Delta N_i = k | \Delta N_i \neq 0)$, $k \in \mathbb{Z}$ is the (integer) jump distribution for technology stage changes.

3.1 Equilibrium Production

We proceed to solve the system (17) in the case where the controls are not coupled, $\mathcal{A}_i(\mathbf{n}) = \mathbb{R}_+^2 \forall \mathbf{n}$. This decouples the optimization problems for q_i, a_i in (17). We also observe that the first term in each supremum captures the instantaneous Cournot competition and contains only q_i , while the other two terms represent the market shifts due to potential innovations and are related to a_i .

Given current cost profile $\mathbf{c}(\mathbf{n})$, equilibrium production is determined as in the static case. Indeed, production rates only affect the immediate profits $\pi_i(\mathbf{c}(\mathbf{n}))$ and hence can be optimized pointwise in time (recall that we restrict to Markov feedback strategies). Thus, $q_i^*(\mathbf{n})$ can be determined from Lemma 1 for the static Cournot oligopoly after re-labeling the players in increasing order of their costs.

In the examples below we focus on the linear price function which has zero relative prudence, $\rho = 0$. In that case, the following Corollary follows immediately from Proposition 3 and Lemma 1.

Corollary 1. *Suppose $P(Q) = \eta - Q$ (i.e. $\rho = 0$ in Corollary 3). Then recalling the notation $B_\ell = \sum_{k=1}^\ell c_k$, we have*

$$q_i^*(\mathbf{c}) = \max(\bar{P} - c_i, 0) \quad \text{and} \quad \pi_i^*(\mathbf{c}) = q_i^*(\mathbf{c})^2, \quad \text{where} \quad \bar{P}(\mathbf{c}) = \min \left\{ \frac{B_\ell + \eta}{\ell + 1} : 1 \leq \ell \leq L \right\}.$$

Moreover, with $\ell^* = \max\{\ell : \eta + B_{\ell-1} - \ell c_\ell > 0\}$, we have

$$\frac{\partial q_i^*}{\partial c_i} = -\frac{\ell^*}{\ell^* + 1} \quad \text{and} \quad \frac{\partial q_i^*}{\partial c_j} = \frac{1}{\ell^* + 1}.$$

Explicitly, the duopoly production rates are given by the solution of the static game with costs $\mathbf{c}(\mathbf{n})$:

$$\bar{q}_i(\mathbf{n}) = q_i^*(\mathbf{c}(\mathbf{n})) = \begin{cases} \frac{\eta + c_j(n_j) - 2c_i(n_i)}{3} & \text{if } \frac{\eta + c_j(n_j)}{2} \geq c_i(n_i) \geq 2c_j(n_j) - \eta; \\ \frac{\eta - c_i(n_i)}{2} & \text{if } \eta + c_i(n_i) < 2c_j(n_j); \\ 0 & \text{if } \eta + c_j(n_j) < 2c_i(n_i). \end{cases} \quad (19)$$

The first case is the duopoly equilibrium, in the second case player i has a monopoly and in the last case player j has a monopoly. Under monopoly, one player has the market all to herself, choosing the monopoly optimal level $q_i^* = (\eta - c_i)/2$. Figure 1 illustrates the resulting duopoly “wedge” as a function of costs $\mathbf{c} = (c_1, c_2)$, (and when $\eta = 1$).

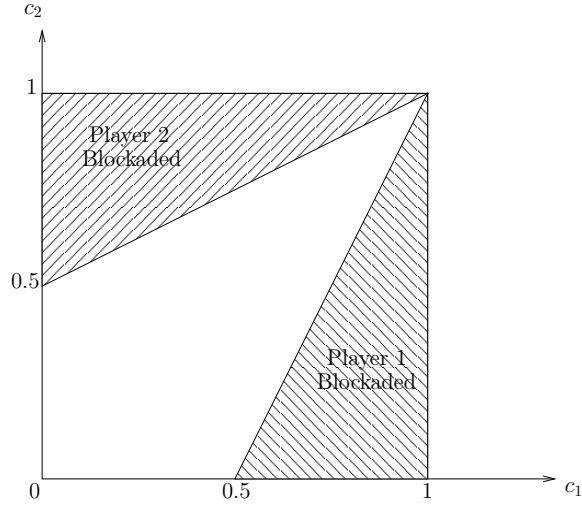


Figure 1: Type of Game Equilibrium in a Cournot Duopoly with linear demand $P(Q) = 1 - Q$.

3.2 Equilibrium R&D Effort

In contrast, the R&D effort level a_i enters (17) in a highly nonlinear fashion. Indeed, a_i appears both in the numerator and denominator of the equations. Fixing a_j^* , the first-order optimality conditions for a_i are obtained by differentiating (17) with respect to a_i . To find the equilibrium level (a_1^*, a_2^*) then requires solving the resulting system of two nonlinear equations.

More precisely, recalling the notation $\pi_i^*(\mathbf{c}(\mathbf{n}))$ for the revenue from production as in (8), and fixing a_j^* , set

$$J_i(a, a_j^*; \mathbf{n}) := \mathbb{E} \left[\frac{1 - e^{-\rho_i \sigma}}{\rho_i} (\pi_i^*(\mathbf{c}(\mathbf{n})) - \mathcal{R}_i(a)) + e^{-\rho_i \sigma} \{ 1_{\{\sigma = \sigma^1\}} v_i(n_1 + 1, n_2) + 1_{\{\sigma = \sigma^2\}} v_i(n_1, n_2 + 1) \} \right]. \quad (20)$$

Then $a_i^*(\mathbf{n}) = \arg \sup_{a \geq 0} J_i(a, a_j^*; \mathbf{n})$. As with production rates, R&D efforts cannot be negative. If the marginal cost of R&D $\mathcal{R}'_i(a)$ is bounded away from zero, it is possible that the equilibrium R&D effort is zero. Indeed, fixing $i = 1$ for concreteness,

$$\frac{\partial}{\partial a_1} J_1(a_1) \Big|_{a_1=0} = \frac{1}{\bar{\lambda} a_2^* + \rho_1} \left[-\mathcal{R}'_1(0) + \bar{\lambda} v_1(n_1 + 1, n_2) - \frac{\pi_1^*(\mathbf{c}(\mathbf{n})) + \bar{\lambda} a_2^* v_1(n_1, n_2 + 1)}{a_2^* + \rho_1 / \bar{\lambda}} \right].$$

Since depending on model parameters $v_1(\mathbf{n})$ and $\pi_1^*(\mathbf{c}(\mathbf{n}))$ can be arbitrarily close to zero, it is clear that if $\mathcal{R}'_1(0) > 0$ then the above expression may be negative and hence $a_1^*(\mathbf{n}) = 0$. In that case, player 1 would invest nothing in R&D and only player 2 innovates, leading to $\mathbb{P}(\sigma = \sigma^2) = 1$. Similarly, it is also possible that $a_1^* = a_2^* = 0$ in which case technologies of both players remain forever frozen. With constant production costs, the corresponding Nash equilibrium for \mathbf{q} reduces to a stationary Cournot game with cost profile $\mathbf{c}(\mathbf{n})$, which is equivalent to the one-shot static market.

Existence of an equilibrium pair (a_1^*, a_2^*) is resolved through the usual method of analyzing the invertibility of the best-response functions, see Grossman and Shapiro [1987], Judd [2003].

3.3 Solving for the Equilibrium Strategies

Once the current equilibrium strategies $a_i^*(n_1, n_2), \bar{q}_i(n_1, n_2)$ are determined, the overall system (17) can be viewed as a double array of nonlinear optimization problems, coupled within each

other and indexed by \mathbf{n} . The standard paradigm of dynamic programming could then be invoked to solve iteratively “backwards” in \mathbf{n} for the v_i ’s. Indeed, given $v_i(n_1 + 1, n_2), v_i(n_1, n_2 + 1)$ one can determine $v_i(n_1, n_2)$ so that starting with some *terminal* conditions $v_i(N_1, \cdot), v_i(\cdot, N_2)$ one may iteratively solve for the full double array of game values. This gives a well-defined numerical recipe for any finite technology ladder \mathcal{C}_i . The next Lemma shows that the finiteness assumption is not too restrictive.

Lemma 2. *Suppose that $\lambda_i = \bar{\lambda}a_i$ and $\mathcal{R}'_i(0) > 0$. Then the dynamic game can be reduced to a finite-stage one, i.e. for n_i ’s large enough for all i , $a_i(\mathbf{n}) = 0$. Conversely, if $\mathcal{R}'_i(0) = 0$ and the ladder \mathcal{C}_i is strictly decreasing then $a_i^*(\mathbf{n}) > 0$ for all \mathbf{n} .*

The proof is given in Appendix B. Because the technology ladders are monotone and $C_i(t) \geq 0$, the production costs live in a bounded interval. Hence, if the number of stages is unbounded, there must be at least one accumulation point \underline{c}_i for the ladder \mathcal{C}_i . Consequently, after sufficiently many innovations, $c_i(n)$ will be very close to \underline{c}_i and the potential for gain becomes arbitrarily small. If the marginal cost of R&D is strictly positive then R&D efforts will strictly dominate any innovation profits and R&D is shutdown. Conversely, if $\mathcal{R}'_i(0) = 0$, then for small enough level of R&D effort, the R&D gains (which are asymptotically linear in a_i) dominate the negligible R&D expenditures.

Lemma 2 highlights the crucial role of $\mathcal{R}'_i(0)$. If $\mathcal{R}'_i(0) = 0$ then $a_i(n) > 0$ for all n , since the marginal cost of effort is negligible for a sufficiently small. Consequently R&D is always employed and the game never ends (unless the technology ladder is finite). On the other hand, if $\mathcal{R}'(0) > 0$, then eventually the marginal cost of R&D strictly dominates any resulting gains and further R&D becomes economically non-feasible. Thus, there must be absorbing game stages, where the players endogenously forgo R&D. We believe the latter situation is both more realistic economically, and also computationally easier, permitting to employ backward recursion to find game values. Games with infinite stages create technical difficulties in defining an equilibrium strategy. Motivated by this discussion, we henceforth focus on the case

$$\mathcal{R}(a) = \frac{1}{2}a^2 + \kappa a, \quad (21)$$

which combines features of the classical quadratic cost structure and the strictly positive marginal cost controlled by the parameter $\kappa = \mathcal{R}'(0)$.

The proof of Lemma 2 implies the following:

Corollary 2. *In the linear duopoly model with the lower bound on costs $\underline{c}_i = 0$, a sufficient condition to guarantee $a_1(\mathbf{n}) = 0$ is*

$$|\pi_1^*(0, c_2(n_2)) - \pi_1^*(c_1(n_1), 0)| \leq \rho^2 \mathcal{R}'_1(0).$$

In particular if $\max(c_1(n_1), c_2(n_2)) \leq \min(\rho_1^2 \mathcal{R}'_1(0), \rho_2^2 \mathcal{R}'_2(0))$ then $\mathbf{a}(\mathbf{n}) = 0$ and the resulting game stage features no R&D.

If $\mathbf{a}(\mathbf{n}) \equiv 0$ then the resulting market is equivalent to a static Cournot game with cost profile $\mathbf{c}(\mathbf{n})$. In particular, we immediately obtain $v_i(\mathbf{n}) = \frac{\pi_i^*(\mathbf{c}(\mathbf{n}))}{\rho_i}$. Taking \mathbf{n} big enough to satisfy the conditions of Lemma 2 allows to use the above as boundary condition for the backward recursion on the lattice $\mathbf{n} \subset \mathbb{N}^2$ for (17). Note that when solving for $\mathbf{a}(\mathbf{n})$ it is possible that $\mathbf{a}(\mathbf{n}) = 0$ emerges as the solution even in an “interior” game stage, which however poses no numerical concerns.

3.4 Game Evolution

For the remainder of the paper we work with a fixed MPE profile (\mathbf{q}, \mathbf{a}) and drop the corresponding $\bar{\cdot}$ and over-bar. The discrete nature of the state \mathbf{N} implies that the global game is partitioned into *stages*, with innovation epochs characterizing transitions among different stages. The piecewise-constant behavior of \mathbf{N} moreover makes each stage equivalent to a stationary Cournot game with state-dependent payoffs. Namely, at each stage players compete in a static Cournot market while also engaging in R&D that will eventually move them to a new stage. Until R&D comes to fruition, players solve the local problem; as soon as there is an innovation by a “winning” player at time $\underline{\sigma}$, a new optimization problem is considered in turn. The fact that the controls in the constructed MPE are a function of $\mathbf{N}(t)$ imply the following

Proposition 5. *The technology state variable \mathbf{N} is a state-dependent Markov chain on \mathbb{N}^2 . Given the current stage (n_1, n_2) , the sojourn time in this state has a memoryless Exponential distribution with mean $\frac{1}{a_1(\mathbf{n})+a_2(\mathbf{n})}$ and the next transition is to $(n_1 + 1, n_2)$ (resp. $(n_1, n_2 + 1)$) with probability $\frac{a_1(\mathbf{n})}{a_1(\mathbf{n})+a_2(\mathbf{n})}$ (resp. $\frac{a_2(\mathbf{n})}{a_1(\mathbf{n})+a_2(\mathbf{n})}$).*

Thus, the global evolution is characterized by patching together the local equilibria described by \mathbf{q}, \mathbf{a} . One can generate time-scenarios of the dynamic Cournot market as follows. Given \mathbf{n} , we draw two independent Exponential r.v.’s $\sigma^i \sim \text{Exp}(\bar{\lambda}a_i(n_1, n_2)), i = 1, 2$ and take $\underline{\sigma} = \sigma^1 \wedge \sigma^2$. We also solve for $\mathbf{q}(\mathbf{n})$ using (19). This yields the game solution on $[0, \underline{\sigma}]$ and based on the relationship between σ^1 and σ^2 the next game stage $\mathbf{N}(\underline{\sigma}) = (n_1 + 1_{\{\sigma^1 < \sigma^2\}}, n_2 + 1_{\{\sigma^2 < \sigma^1\}})$. After transitioning, the new production rates $\mathbf{q}(\mathbf{N}(\underline{\sigma}))$ are updated and the above process is repeated with fresh draws for σ^i . By Lemma 2, this chain will reach an absorbing state (where $\mathbf{a} = 0$) with probability 1, so that all other states are transient and the above algorithm is guaranteed to terminate.

Remark 4. The local structure of the Cournot games that arise allows further indexing of game parameters by \mathbf{n} . For example, for a realistic calibration it would be reasonable to assume that the R&D costs depend on the current technology, i.e. $\mathcal{R}(a) = \mathcal{R}(a; \mathbf{n})$. Similarly, one can also index the discount factors or the demand curves by \mathbf{n} and retain the Markov structure for dynamic programming.

Remark 5. We have focused in this section on the duopoly competition ($L = 2$) to illustrate the effects of competition in a presentable way. Two extreme cases are also of interest: the monopoly case $L = 1$ and the limit $L = \infty$. In the next section, we analyze the case where only one player can innovate, but interacts with a second player through Cournot competition. This is contrasted with the case of two innovators in Section 5, specifically the discussion at the end of that section. This gives some insight into the effect of the number of players that are involved in R&D. In adding more players, choices have to be made as to what their costs of production are if comparisons are to be made, and perhaps only the symmetric situation of adding more identical players makes sense for comparison, where here we are very much interested in heterogeneous situations, even if the players differ only in their initial costs. Analyzing a dynamic game where players may enter over time as their costs drop is interesting too, but requires a separate work. For instance, a case without innovation where players enter as others use up an exhaustible resource is studied in Ledvina and Sircar [2012].

The limit $L \rightarrow \infty$ of a large number of players relates to much current research in (continuum) mean field games. These can be tractable in some situations, but again require a separate and mathematically involved analysis which is beyond the scope here. We refer to Chan and Sircar [2014], Chan and Sircar [2015] and Guéant et al. [2011] for recent developments in mean field games and Cournot competition.

4 Unilateral R&D

We begin our illustrations with a toy model wherein only a single player (for concreteness P1) can engage in R&D. This can be viewed as a special case where the second player is constrained to $a_2 \equiv 0$. Since $C_2(t) \equiv c_2$ is now fixed, the stochastic state is the one-dimensional ($N_1(t)$) and the system (17) reduces to a coupled array of optimization problems for vectors rather than matrices. Unilateral R&D can be visualized as a horizontal movement in the diagram of Figure 1, from right to left at the given level c_2 . Note that interpreting equilibrium earnings of player 1 as a generic technology-dependent reward function, the model with a single innovator is abstractly equivalent to incorporating intermediate rewards into the multi-stage racing game studied by Judd [2003].

For the remainder of this section we fix c_2 and drop it from arguments of all relevant functions. Also for concreteness we take the linear inverse demand curve $P(Q) = \eta - Q$ from Corollary 1, yielding

$$\pi_i^*(\mathbf{c}(n)) = (q_i^*((c_1(n), c_2)))^2.$$

By Lemma 2, without loss of generality the technology of P1 has a finite number of stages \bar{n} . At the ultimate stage, no more R&D research is possible for either player and we face the static-cost Cournot market with game values

$$v_1(\bar{n}) = \frac{\pi_1^*(c_1(\bar{n}), c_2)}{\rho_1}, \quad v_2(\bar{n}) = \frac{\pi_2^*(c_1(\bar{n}), c_2)}{\rho_2}.$$

The rest of the game values are determined via the system of difference equations (obtained by setting $a_2 \equiv 0$ in (17))

$$v_1(n) = \sup_{a \geq 0} \frac{1}{\lambda a + \rho_1} \{ \pi_1^*(c_1(n), c_2) + \bar{\lambda} a v_1(n+1) - \mathcal{R}_1(a) \}. \quad (22)$$

The above is a nonlinear equation for $a_1(n)$ with $v_1(n+1)$ entering as a coefficient. For instance, assuming quadratic costs of (21), $a_1(n)$ is the root of the quadratic from the first-order-condition equation

$$-\rho_1 \bar{\lambda} v_1(n+1) + (a + \kappa)(\bar{\lambda} a + \rho_1) + \bar{\lambda}(\pi_1^*(c_1(n), c_2) - (a^2/2 + \kappa a)) = 0.$$

The game values for player 2 are similarly determined from

$$v_2(n) = \frac{1}{\bar{\lambda} a_1(n) + \rho_2} \{ \pi_2^*(c_1(n), c_2)^2 + \bar{\lambda} a_1(n) v_2(n+1) \}.$$

Globally, (22) can give rise to three market structures based on the cases in (19): P1 monopoly, P2 monopoly, and duopoly. In particular, for $\eta/2 < c_2 < \eta$, as player 1 innovates all three structures arise dynamically. Indeed, as C_1 is reduced, the market moves from old-generation monopoly by player 2, to duopoly, and finally to new-generation monopoly of player 1. The phase transitions take place at $C_1 = \frac{\eta+c_2}{2}$ and $C_1 = 2c_2 - \eta$, respectively, cf. Figure 1. (If $c_2 < \eta/2$, P1 never achieves monopoly.)

Figure 2 illustrates this solution in the case of a linear technology ladder $c_1(n) = 1 - \mu n$ and $P(Q) = 1 - Q$. We show three investment curves $n \mapsto a_1(n)$ for different costs of the other player c_2 . In scenario (a), $c_2 = 0.7$ which represents weak competition; in particular once $c_1(n) \leq 0.4$, P2 is blockaded and P1 has monopoly. In scenario (b), $c_2 = 0.4$ is proxy for moderate competition; scenario (c) $c_2 = 0$ illustrates strong competition since $c_2 \leq c_1(n)$ for all n . In both of the latter cases P2 always produces and is never blockaded. In all three cases we also have that for $c_1(n)$ large (n small), P1 is blockaded, $q_1^*(c_1(n), c_2) = 0$. These stages are illustrated with filled symbols in Figure 2 (for example in case (a), $q_1^*(c_1(n), c_2) = 0$ for $n \leq 7$).

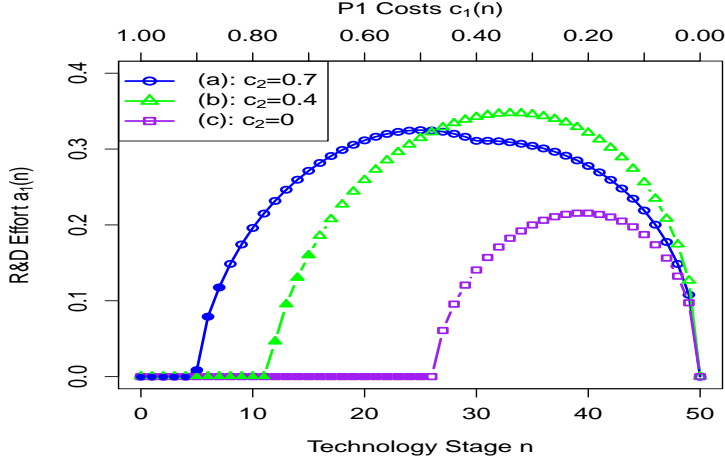


Figure 2: Effort curves for unilateral R&D in a Cournot duopoly. The technology ladder is $c_1(n) = 1 - 0.02n$, $n = 0, \dots, 50$ with effort costs $\mathcal{R}(a) = a^2/2 + 0.3a$, $\rho_1 = 0.01$ and $P(Q) = 1 - Q$. Filled symbols indicate stages where $q_1^*(c_1(n), c_2) = 0$.

4.1 Value of R&D

With unilateral R&D the Markov chain N is a birth process (Markov counting process), and the game evolution simply proceeds sequentially through stages $(c_1(1), c_2), (c_1(2), c_2), \dots$. Denote by $0 = \tau^0, \tau^1, \tau^2, \dots$ the respective transition times to the next technology stage. The τ^k 's are determined by the effort levels $a_1(n)$, namely $\tau^n - \tau^{n-1} \sim \text{Exp}(\bar{\lambda}a_1(n))$, so that $\mathbb{E}[\tau^n - \tau^{n-1}] = (\bar{\lambda}a_n)^{-1}$. We can express the value of player 1 via

$$\begin{aligned}
v_1(1) &= \mathbb{E}_0 \left[\int_0^{\tau^1} e^{-\rho_1 s} (\pi_1^*(c_1(1), c_2) - \mathcal{R}(a_1(1))) ds + \int_{\tau^1}^{\tau^2} e^{-\rho_1 s} (\pi_1^*(c_1(2), c_2) - \mathcal{R}(a_1(2))) ds \right. \\
&\quad \left. + \int_{\tau^2}^{\tau^3} e^{-\rho_1 s} (\pi_1^*(c_1(3), c_2) - \mathcal{R}(a_1(3))) ds + \dots + \int_{\tau^{\bar{n}}}^{\infty} e^{-\rho_1 s} \pi_1^*(c_1(\bar{n}), c_2) ds \right] \\
&= \left\{ \sum_{n=1}^{\bar{n}-1} \frac{(\pi_1^*(c_1(n), c_2) - \mathcal{R}(a_1(n)))}{\rho_1} \mathbb{E}_0 \left[e^{-\rho_1 \tau^{n-1}} - e^{-\rho_1 \tau^n} \right] \right\} + \frac{\pi_1^*(c_1(\bar{n}), c_2)}{\rho_1} \mathbb{E}_0 [e^{-\rho_1 \tau^{\bar{n}}}] \\
&= \sum_{n=1}^{\bar{n}+1} \{ \pi_1^*(c_1(n), c_2) - \mathcal{R}(a_1(n)) \} \cdot \left\{ \frac{1}{\bar{\lambda}a_1(n) + \rho_1} \prod_{j=1}^{n-1} \left(\frac{\bar{\lambda}a_1(j)}{\bar{\lambda}a_1(j) + \rho_1} \right) \right\}, \tag{23}
\end{aligned}$$

where the last equality is based on the expansion $e^{-\rho_1 \tau^n} = e^{-\rho_1 \tau^1} e^{-\rho_1 (\tau^2 - \tau^1)} \dots e^{-\rho_1 (\tau^n - \tau^{n-1})}$ and the moment generating function of Exponential random variables. The above formula links the stagewise payoff rates $\pi_1^*(c_1(n), c_2)$ and R&D rates $a_1(n)$ to the game value of player 1.

This view also highlights the value of R&D. Compared to the base scenario of zero R&D investment, the *gap*

$$G_1(n) = v_1(n) - \int_0^{\infty} e^{-\rho_1 s} \pi_1^*(c_1(n), c_2) ds = v_1(n) - \frac{\pi_1^*(c_1(n), c_2)}{\rho_1} \tag{24}$$

precisely captures the NPV of R&D investments. Note that $G_1(n) = 0$ is equivalent to $a_1(n) = 0$ and represents the end of the technology ladder. We find, cf. right panel of Figure 3, that the

structure of $n \mapsto G_1(n)$ is closely related to the investment curve $n \mapsto a_1(n)$ and has an inverted-U shape. For n small, R&D gains are small since current revenues are still low (i.e. discounting makes the NPV minimal). For n large, all the technology gains are nearly exhausted so necessarily $G_1(n)$ is also small. Thus, $G_1(n)$ is maximized in the middle stages which is also when the R&D effort is strongest.

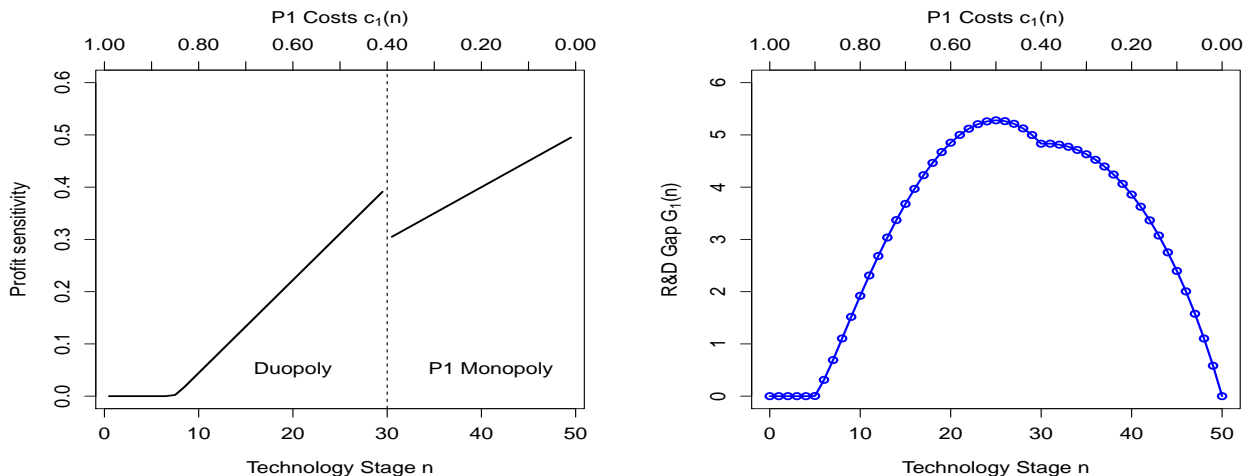


Figure 3: Left: sensitivity of instantaneous profit to production costs $-\partial\pi(c_1, c_2)/\partial c_1$. P2 costs are fixed at $c_2 = 0.7$. Right: gap $G_1(n) = v_1(n) - \frac{\pi_1(n)}{\rho_1}$ between the NPV from optimal R&D and NPV from zero R&D. Linear monopoly model with parameters as in Figure 2.

4.2 Dynamic Market Structure

In this section we investigate how the two facets of our model – R&D efforts a_1^* and production rates \bar{q}_1 interact across game stages, in particular in terms of the resulting Cournot market structure. The dependence of $a_1^*(n)$ on n can be inferred from (23). For simplicity, assume the linear ladder $c_1(n) = 1 - \mu n$. In that case, the immediate gain from a technology advance is proportional to $\partial\pi_1^*/\partial c_1$ which, by Corollary 1, is

$$\left. \frac{\partial\pi_1^*}{\partial c_1} \right|_{c_1=c_1(n)} = -2 \frac{\ell(n)}{\ell(n) + 1} q_1^*(c_1(n), c_2). \quad (25)$$

While production q_1^* is increasing in n , the number of active players $\ell(n)$ switches midway from duopoly to a monopoly, generating a downward jump in $\partial\pi_1^*/\partial c_1$ (cf. left panel of Figure 3). This occurs because a monopolist is less sensitive to lowering costs than a duopoly player. Consequently, the link between n and $a_1^*(n)$ is driven by two competing effects: the finite horizon effect (costs cannot be decreased below \underline{c}) and the cost-sensitivity effect (equation (25)).

When C_1 is large (n small), P1 is the outsider and the low revenues lead to low R&D investment. Moreover, the cost-sensitivity effect dominates (since low costs are yet too far into the future) and $a_1^*(n)$ is increasing in n . As C_1 decreases (n increases), P1 catches up to P2 in terms of production costs and R&D investment rises, peaking roughly when $c_1(n) = c_2$. In the end stages (n large), $c_1(n)$ is already very low and the potential for future technology gains diminishes. As a result, the horizon effect dominates, causing R&D investment to slow down and making $a_1^*(n)$ downward sloping. Ultimately, $a_1^*(\bar{n}) = 0$ and the game terminates. Thus, over time, R&D investment goes through a boom-and-bust or inverted-U pattern. Such investment patterns were observed as

early as Kamien and Schwartz [1978], though in their work it was due to the exhaustibility of the underlying production resource, whose depletion eventually makes R&D worthless.

Next, we discuss the impact of R&D on the production *rates*. As indicated by Lemma 1, lower production costs C_1 increase \bar{q}_1 and Q . R&D naturally crowds out player 2 whose production \bar{q}_2 declines, creating a substitution effect whereby increased production from player 1 is partially offset by decreased output of player 2. It follows that the price $P_t = \eta - \bar{q}_1(N(t)) - \bar{q}_2(N(t))$ is non-increasing over time. Moreover, this price decline accelerates once player 1 achieves monopoly, since in monopoly the above substitution effect goes away.

The inverted U-shape of the investment curve is broadly similar across all three scenarios in Figure 2, though several variations appear. Notably, scenario (a) features a double boom-and-bust pattern. This occurs due to a structural shift in the market as P1 achieves monopoly for $n \geq 0$. The different sensitivity to production costs (monopolist having smaller cost-convexity $\partial^2 \pi_1^* / \partial c_1^2$) between a monopolist and a duopolist causes the R&D investment to slightly increase again around $n = 30$ after falling for several stages. We also note that $c_2 \mapsto a_1^*(n; c_2)$ is not monotone— e.g., the investment curves for $c_2 = 0.7$ and $c_2 = 0.4$ cross. This illustrates the non-monotone relationship between market competitiveness and R&D investment – while very strong competitors (scenario (c)) discourage R&D by diminishing future gains, moderate competition can actually spur R&D as producers try to get an edge. Moreover, the peak of R&D efforts depends strongly on c_2 – in scenario (b) R&D is maximized around stage $n = 35$, while in scenario (a) it is peaking much earlier around $n = 25$.

Figure 2 also demonstrates the nontrivial interaction between market structure and R&D. Recall that either $\bar{q}_1(n) = 0$ (blockading) or $a_1^*(n) = 0$ (technological stagnation) are possible, leading to four regimes for the production + R&D controls. All four cases may be observed in Figure 2. First, the case $\bar{q}_1(n) = 0$ takes place when $c_1(n)$ is too big and P1 is blockaded. In that case, P1 receives no current revenues and hence any respective R&D investment is solely about hoping for future profits. Those may or may not be enough to justify current effort. To wit, in scenario (a) for $n \in \{5, 6, 7\}$ we have $a_1^*(n) > 0$ even with $\bar{q}_1(n) = 0$. This is the proverbial “invest now for the better future” situation. Simultaneously in the same scenario when $n \leq 4$, $a_1^*(n) = 0 = \bar{q}_1(n)$ – the “light in the end of the tunnel” is too far away so that R&D is not undertaken. Indeed, with $n = 4$ one needs at least 4 technology improvements to generate any revenue; the associated revenues are so far into the future that they are not sufficient to finance the immediate R&D expenditures, cf. (23). As a result, P1 invests nothing in R&D and will remain forever blockaded, leading to $v_1(n) = 0$ for $n \leq 4$. In contrast for $n \geq 5$, only 3 or fewer technology improvements are necessary to break P2’s monopoly and enter the market and R&D becomes economically worthwhile. A similar situation happens in scenario (b). However, in scenario (c) we see the opposite effect— at stage $n = 27$ and $c_1(n) = 0.48$, $\bar{q}_1(n) > 0$ but $a_1^*(n) = 0$. In other words, even with positive present cashflows, the expected gain from technology investment is not big enough to justify it. This happens because P2 is too strong so P1 will continue to have a small market share for the foreseeable future. As a result, P1 “gives up” on R&D despite being in the market.

4.3 Policy Implications

The above observations could be translated into policy recommendations. We imagine that P2 represents an entrenched incumbent (for example fossil fuel energy generators) and P1 is a new-technology entrant that policy makers wish to support (eg. a new method for green energy production). In an ideal world, the future benefits of P1’s tools would allow to privately finance the requisite R&D even before the new method is generating any revenues. However, if these gains are too remote, policy intervention via R&D subsidies might be necessary to get a head start. For example, starting at $n = 0$ and scenario (a), subsidized R&D is needed for the first 4 stages otherwise

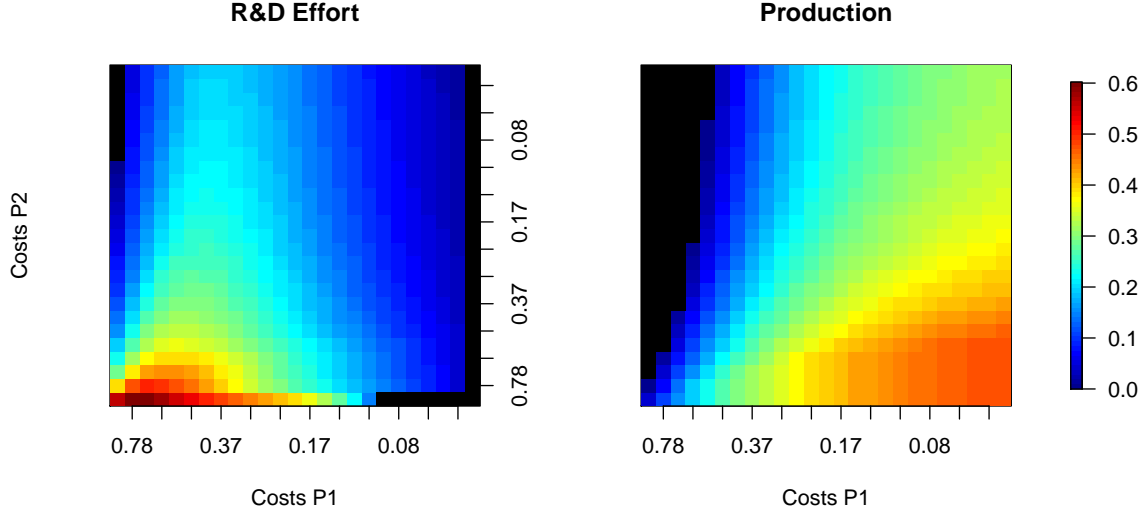


Figure 4: Duopoly R&D race with linear price function $P(Q) = 1 - Q$. Right panel shows the effort $a_1^*(n_1, n_2)$ and left panel the production rate $\bar{q}_1(n_1, n_2)$. Quadratic costs $\mathcal{R}_i(a) = a^2/2 + 0.1a$ with $\lambda_i(t) = 5a_i(t)$, $\rho = 0.1$. $c_i(n) = e^{-n/8}$, $\mathbf{n} \in \{1, \dots, 25\}^2$.

the new technology will never be developed on its own. However, once $n \geq 5$, the subsidy can be withdrawn and replaced with private financing. One could imagine various tax/equity mechanisms that recoup these seed investments in the ultimate future $\bar{n} = 50$. Scenario (c) shows that such kick-start subsidies could be needed even if P1 is already in the market but the initial R&D hurdle is too high (the situation where $a_1^*(n) = 0, \bar{q}_1 > 0$). Overall, Figure 2 supports the idea of seed subsidies that can then be withdrawn once the new technology is sufficiently competitive. Subsidies are especially attractive if one imagines the government to have lower discount rates compared to the game players, facilitating long-maturity loans.

5 Bilateral R&D Race

To illustrate a truly dynamic R&D race, we next investigate a duopoly $L = 2$ with bilateral R&D strategies $a_1(t), a_2(t)$. We continue to maintain linear inverse demand $P(Q) = \eta - Q$. For expositional clarity we focus on the symmetric case where the R&D costs $\mathcal{R}(a)$, discount factors $\rho_i \equiv \rho$, and technological ladders \mathcal{C}_i of both players are identical. Of course, during the evolution of the game, the players will end up in different stages, making the sub-games non-symmetric. In terms of the diagram in Figure 1, the games moves from upper-right to lower-left by taking steps either to the South or West.

Figure 4 shows the optimal feedback controls for a geometric technology ladder $c_i(n) = \exp(-\mu n)$ and quadratic costs (21). Since absolute progress slows down, there is an economically optimal limit level \bar{n} such that no R&D takes place beyond stage (\bar{n}, \bar{n}) . Based on Lemma 2, we can take $\bar{n} = -\mu^{-1} \log(\rho^2 \kappa)$. After reducing to a finite-stage setting, we can inductively solve for $v_i(\mathbf{n})$ over the resulting square lattice $n_i \in \{1, \dots, \bar{n}\}, i = 1, 2$.

The two panels of Figure 4 show the feedback controls $\bar{q}_1(\mathbf{n}), a_1^*(\mathbf{n})$ of player 1 as a function of \mathbf{n} . By symmetry, the solution for P2 is simply the mirror image $\bar{q}_2(n_1, n_2) = \bar{q}_1(n_2, n_1)$, etc. In the right panel we plot the equilibrium production rate $\bar{q}_1(n_1, n_2)$ which as expected is maximized

when n_1 is large and n_2 is small, i.e. $C_1 \ll C_2$. In the first few stages ($n_1 < 5$), it is possible that P1 is blockaded ($\bar{q}_1(n_1, \cdot) = 0$ in the upper-left corner). In the left panel, we show the corresponding effort level $a_1^*(n_1, n_2)$. For $n_1 = 25 = \bar{n}$ (right boundary), $a_1^*(\bar{n}, \cdot) \equiv 0$ no more R&D is undertaken. Similarly, no R&D is undertaken when $n_1 \gg n_2$ (upper-left corner) in which situation P1 is too “behind” P2 and is moreover blockaded, making eventual R&D gains too minuscule to be feasible. However, we note that the region $\{(n_1, n_2) : a_1^*(n_1, n_2) = 0\}$ is much smaller than the blockading region for P1, so that in many game scenarios P1 may be blockaded but innovating, hoping to catch up to P2 eventually.

Most interestingly, Figure 4 demonstrates that the R&D investment is maximized when n_1 is slightly larger than n_2 . This is the situation where P1 has a small technological advantage over P2, whereby she has the highest motivation (and the funds) to increase the gap with her competitor. In that sense, a symmetric competition $N_1(t) = N_2(t)$ is unstable, since whichever player is ahead, is also putting more in R&D than the player that is behind. This feature confirms Lemma 1 which shows that *ceteris paribus* a more dominant player is more sensitive to her costs and hence more incentivized to innovate.

Over time, the R&D race reaches one of the absorbing stages where $a_i^*(n_1, n_2) = 0 \forall i$. In the present example this effectively reduces to $\lim_{t \rightarrow \infty} \mathbf{N}(t) = (\bar{n}, \bar{n})$ (there is also a small absorbing region in the extreme case of $n_1 \gg n_2$ which is however extremely unlikely starting from $N_1(0) = N_2(0)$). Figure 5 shows the distribution of $(N_1(t), N_2(t))$ over $t = 2, 8, 15, 25$ based on the initial condition $\mathbf{N}(0) = (1, 1)$. We observe that the duopoly instability causes a rather large spread (or “variance”) at intermediate times $t = 8, 15$ where a significant number of sub-stages could be realized. As $t \rightarrow \infty$, the game collapses back to the main attractor in the upper-right corner.




Figure 5: Distribution of $(N_1(t), N_2(t))$ in \mathbb{N}^2 in a duopoly R&D race. All parameters are as in Figure 4. Each panel shows a 2-D histogram of (N_t^1, N_t^2) based on 500 simulated paths; size of each gridpoint is proportional to the empirical frequency. The initial condition was $\mathbf{N}(0) = (1, 1)$.

Comparing back to the setting of Section 4, the difference between the unilateral and bilateral settings sheds light on the impact of R&D competition on the competitive behavior of the producers. From the point of view of P1, if she has the sole access to innovation and engages in unilateral R&D then she is able to fully reap the resulting benefits; in contrast bilateral R&D competition necessarily dilutes gains from lower costs as P2 is expected to catch up over time and maintain competitive pressures. This intuition suggests that the P1 game value ought to be larger under unilateral R&D than under bilateral R&D, and consequently that her effort should be lower (since the NPV of R&D is generally less). While the former statement is true, the latter is not. To wit, in our numerical experiments we found that in some scenarios, fixing P2 production costs at a given level c_2 and comparing the resulting $a_1^{Un}(n)$ (from Section 4) against $a_1^{Bi}(n)$ (from this Section), the

bilateral-R&D effort is *higher*. In other words, removing R&D competition makes immediate R&D research less urgent for P1, leading to lower optimal effort. This finding could be connected to the concept of pre-emption, whereby bilateral R&D forces producers to exert additional effort in order to stay ahead (or not fall behind) competitors, which in turn further erodes their profitability.

6 Effect of Uncertainty

In the model (4) the only stochastic source are the counting processes N_i controlling technological advances. It is of interest to understand the impact of the discrete shocks in N_i on the game structure. In particular, a natural question is whether uncertainty impedes or encourages R&D. To shed some insight, we consider *refinements* of the technology ladder \mathcal{C} . Namely, starting with a ladder \mathcal{C} , consider a refined ladder \mathcal{C}' which satisfies

$$c'(2n) = c(n) \quad \text{and} \quad c'(2n+1) = \frac{c(n) + c(n+1)}{2}.$$

To compensate for having more stages, we also double the scaling constant $\bar{\lambda}' = 2\bar{\lambda}$. Informally, the refined model represents the situation where R&D discoveries are half the size but occur twice as frequently (for same amount of effort). Let $(N'(t))$ be the counting process in the refined model. With this setup and keeping $a_i(t)$ fixed, we observe that

$$\mathbb{E}[N_i(t)] = \mathbb{E}[N'_i(t)/2], \quad \text{Var}(N_i(t)) = 2\text{Var}(N'_i(t)/2).$$

Therefore, using $N'/2$ as the measure of technology progress in the refined model, the refined model has same rate of R&D discoveries, but only half of the corresponding variance.

As the degree of refinement is increased, the variance of R&D innovations converges to zero. Thus, the above construction offers a way to interpolate between the deterministic model (zero variance) and the stochastic version analyzed so far. To be precise, consider a continuous decreasing map $f(x) \in [0, 1], x \in [0, \bar{x}]$, where x represents the continuous level of “progress”. Innovation may go on indefinitely, or have some finite bound \bar{x} . The map $x \mapsto f(x)$ defines the technology ladder, e.g. $f(x) = \exp(-\mu x)$. Now for any $M \in \mathbb{Z}$, let $c(n; M) = f(n/M)$, $n = 1, \dots$, which gives rise to a ladder $\mathcal{C}^{(M)}$ with a corresponding scaling $M\bar{\lambda}$. We refer to M as the refinement level. The case of finite M nests the discrete technology ladders while in the limit $M \rightarrow \infty$, one obtains a deterministic model where progress is completely dependent on the R&D level without any intrinsic uncertainty. For ease of comparison across different refinement levels, for the remainder of this Section we re-parameterize both q and a in terms of x , $q(x; M) = q(c(n; M))$ keeping in mind that $x \in \{n/M : n = 1, \dots, \}$.

The impact of uncertainty on the game equilibrium can now be analyzed by the dependence of the equilibrium strategies on M . Note that $q(x) = q(f(x))$ only depends on the present costs and hence is independent of M ; in contrast $a(x)$ is driven by the amount of uncertainty corresponding to the refinement level. As M rises, the corresponding dynamic game has more and more stages, so that players have more opportunities to adjust their feedback-form controls and therefore extract more value. However, uncertainty also benefits the players because it allows for potentially rapid progress which is assigned extra weight due to the discounting involved. Indeed, the map $\tau \mapsto \exp(-\rho\tau)$ is convex, so that a mean-preserving transformation of τ will affect the corresponding expectation (cf. (23)) inversely in terms of its variance. The next Lemma shows that the overall impact of M on game values is as a result ambiguous.

Lemma 3. *Consider a technology ladder $\mathcal{C}^{(M)}$ and a doubled ladder $\mathcal{C}^{(2M)}$. Then $v^{(2M)}(x)$ may be either bigger or smaller than $v^{(M)}(x)$.*




Figure 6: Comparison of game values $v_1(\cdot, n_2)$ and effort levels $a_1^*(\cdot, n_2)$. Bilateral symmetric R&D game with $\rho = 0.1$ and $\mathcal{R}(a) = a^2 + 0.2a$ and linear technology progress $c_1(n) = 1 - n/M$, $\bar{\lambda} = 0.4M$ for $M = 10, 20, 50$. We show the results for P1 at the fixed level $c_2 = 0.7$ of P2 costs.

Put another way, less uncertainty lowers the probability of extra profits from rapid progress but allows more flexibility in R&D investments over time. The above results should be contrasted with the findings of Lafforgue [2008] in the single-agent setting. Lafforgue finds that uncertainty unambiguously lowers innovation (and therefore overall utility). This is because the stationary solution obtained in Lafforgue [2008] eliminates any extra potential gains from favorable innovation scenarios. In contrast, in our model the additional gains that arise from for example the duopoly-to-monopoly transition, percolate through the game stages and make the unlikely but potentially very lucrative rapid innovation paths strategically important. Thus, increased uncertainty can sometimes benefit the player and make her even more enthusiastic about R&D. (The non-monotone R&D effort in Figure 2 is again instructive here.)

Figure 6 shows the impact of R&D success volatility in a symmetric bilateral model. We take the linear ladder $f(x) = 1 - \mu x$, $x \in [0, 1/\mu]$; in this case refinement is equivalent to taking $c_1(n; M) = 1 - n/M$, $n = 1, \dots, M$ and varying M . To illustrate Lemma 3 we observe that as M changes, both R&D efforts $a_i(x)$ and game values $v_i(x)$ have ambiguous changes, sometimes increasing and sometimes decreasing.

6.1 Deterministic Limit

The deterministic case is obtained by taking the formal limit $M \rightarrow \infty$, which corresponds to $\mathcal{C}^{(M)} \rightarrow f$. The fluid limit of $N_i(t)/M$ yields that the underlying technology state x slides down the ladder $x \mapsto f(x)$ at the speed proportional to $a_i(t)$, i.e. $dx_i(t) = a_i(t) dt$ or equivalently for $C_i(t) = f(x_i(t))$

$$dC_i(t) = f'(f^{-1}(C_i(t))) \cdot a_i(t) dt.$$

For presentational clarity, we revisit the deterministic version of the unilateral R&D model, where only P1 innovates and the technology ladder is linear $f(x) = 1 - \bar{\lambda}x$, so that

$$dC_1(t)/dt = -\bar{\lambda}a(t).$$

Fix c_2 and denote by $\tilde{g}_i(c)$ the value function of player i starting with initial P1 costs $C_1(0) = c$. To distinguish from the stochastic model, the corresponding controls are denoted $\tilde{a}(c)$ and $\tilde{q}(c)$.

Taking quadratic R&D costs (21), \tilde{g}_1 solves the nonlinear ODE

$$\rho_1 \tilde{g}_1(c) = \pi_1^*(c, c_2) + \frac{(-\bar{\lambda} \tilde{g}'_1(c) - \kappa)_+^2}{2} = 0, \quad c \in [0, 1], \quad (26)$$

where $\pi_1^*(c, c_2)$ is the equilibrium profit under cost profile (c, c_2) (see Lemma 1) and the initial condition is $\tilde{g}(0) = \frac{\pi_1^*(0, c_2)}{\rho_1}$. Note that \tilde{g}_1 is decreasing, hence the minus sign in front of \tilde{g}'_1 . The optimal controls are $\tilde{q}(c) = q_1^*(c, c_2)$ from (19) and $\tilde{a}(c) = (-\bar{\lambda} \tilde{g}'_1(c) - \kappa)_+$. The relationship in (26) subsumes the static case $\tilde{a}(c) = \tilde{q}(c) = 0$ whereby $\tilde{g}_1(c) = 0$, and also the case of zero R&D that implies $\tilde{g}_1(c) = \rho_1^{-1} \pi_1^*(c, c_2)$ or in differential form $\rho_1 \tilde{g}'_1(c) = \frac{\partial}{\partial c} \pi_1^*(c, c_2)$.

Note that at $c = 0$, $\frac{(-\lambda \tilde{g}'_1(0) - \kappa)_+^2}{2} = \rho_1 \tilde{g}_1(0) - \pi_1^*(0, c_2) = 0$, so the boundary condition is singular ($\tilde{g}'(0)$ is undetermined). For a numerical solution we instead use $\lim_{c \downarrow 0} \tilde{g}'_1(c) = -\kappa/\bar{\lambda}$ which corresponds to the condition that R&D efforts are smoothly wound down to zero as the ultimate technology state is approached. As before, one can view $\rho_1^{-1} \pi_1^*(c, c_2) \leq \tilde{g}_1(c)$ as an ‘‘obstacle’’ for \tilde{g}_1 , with $\rho_1 \tilde{g}_1(c) = \pi_1^*(c, c_2)$ if and only if $\tilde{a}(c) = 0$. The gap $G(c) := \tilde{g}_1(c) - \rho_1^{-1} \pi_1^*(c, c_2)$ represents the NPV of R&D gains and substituting into (26) is shown to have a quadratic relationship to the R&D effort level $\tilde{a}(c)$,

$$G(c) = \frac{1}{2\rho_1} \tilde{a}(c)^2. \quad (27)$$

Relationship (27) sheds light on the link between the right panel of Figure 3 and the effort curves in Figure 2.

Moreover, recall that the profit function $\pi_1^*(c, c_2)$ itself is piecewisely defined in terms of the three possible market structures (duopoly or monopoly of either producer). As a result, the full solution for $\tilde{g}_1(c)$ is defined piecewise over as many as four intervals, with both pre-specified and free boundaries arising. The ODE (26) clarifies the phase transitions resulting from changes in market structure. Indeed, formally differentiating (26) we obtain

$$\rho_1 \tilde{g}'_1(c) = \frac{\partial}{\partial c} \pi_1^*(c, c_2) - \bar{\lambda} \tilde{g}''_1(c) \cdot (-\bar{\lambda} \tilde{g}'_1(c) - \kappa)_+.$$

We now observe that since $c \mapsto \pi_1^*(c, c_2)$ is continuous, by (26) so is $c \mapsto \tilde{g}'(c)$ (and hence $c \mapsto \tilde{a}(c)$). At the same time, at levels where the market changes from a monopoly to a duopoly and vice-versa, the cost-sensitivity of profit changes. This implies that $\tilde{g}''(c)$ must be discontinuous, introducing higher-order discontinuities into \tilde{g}_1 . In particular, because a monopolist is less sensitive to her costs, at the transition level $c = 2c_2 - 1$, we have

$$-\frac{\partial}{\partial c} \pi_1^*(c, c_2) \Big|_{c=(2c_2-1)-} < -\frac{\partial}{\partial c} \pi_1^*(c, c_2) \Big|_{c=(2c_2-1)+},$$

and this jump may cause the sign of $\tilde{g}''(c)$ to change from negative to positive when transitioning to P1 monopoly. In other words, the \tilde{g} may go from concave to convex, and \tilde{a} may go from decreasing to increasing, generating a kink at $c = 2c_2 - 1$.

Figure 7 illustrates these effects for the case $c_2 = 0.6$. In that case, there are two fixed phase boundaries at $c = 0.8$ and $c = 0.2$ which indicate blockading of P1 and P2 respectively. The Figure confirms the mentioned double inverted-U relationship for $c \mapsto \tilde{a}(c)$. In particular, there is a kink at $c \simeq 0.45$ (where $-\frac{\partial}{\partial c} \pi_1^*(c, c_2) = \rho_1 \kappa/\bar{\lambda}$) which causes R&D to be started with a ‘‘jolt’’, and another kink at $c = 0.2 = 2c_2 - 1$ which causes R&D trend to suddenly change from negative to positive. Economically, the time-trend of current R&D investments is tied to the change in expectations of

future gains from lowered costs. Thus, as $C_1(t)$ drops below 0.3, the gains to P1 from R&D diminish as she anticipates lower cost-convexity (the downward jump in $\frac{\partial}{\partial c}\pi_1^*(c, c_2)$ observed in Figure 3) once in monopoly. As a result, R&D effort is lowered. However, once monopoly is actually achieved, $t \mapsto \partial\pi_1^*(C_t, c_2)/\partial c$ resumes its growth and the R&D trend is reversed.

Figure 7 also compares the deterministic R&D effort curve $\tilde{a}(c)$ to the discrete versions $a^{(M)}(c)$ obtained with a finite M . We observe that while for $M = 10$ the discretization makes a significant difference, already with $M = 100$ stages, the stochastic setting is very close to the deterministic one. Notably, for $c = 0.4$ and $M = 10$, the large uncertainty (and hence potential for quick gains) makes it optimal to do some R&D, while for $M = 100$ and the deterministic case, at the same production cost level c the expected gains are not sufficient to justify investments. This highlights the extra benefits that uncertainty may offer and shows that correct modeling of uncertainty is an important aspect of the framework.

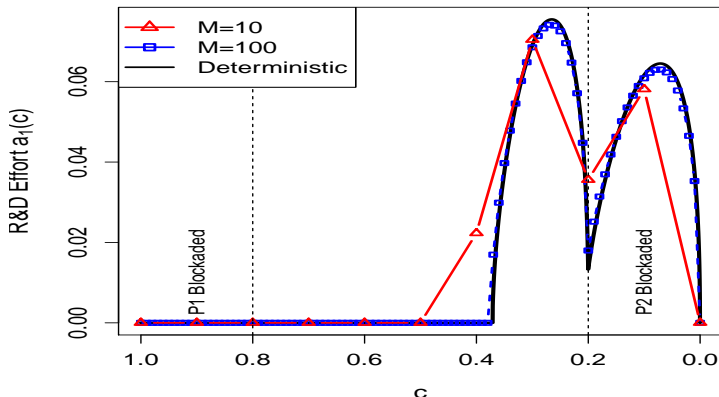


Figure 7: Effect of uncertainty on effort curves. The deterministic technology ladder is $c_1(x) = f(x) = 1 - \bar{\lambda}x$ with $\bar{\lambda} = 0.01$. We also show the discretized cases $M = 10, 100$ where $c_1(n) = 1 - n/M$ and discovery rate is $M\bar{\lambda}$. Effort costs $\mathcal{R}_1(a) = a^2/2 + 0.4a$, $\rho_1 = 0.01$ and $c_2 = 0.6$.

Remark 6. A related deterministic Cournot *duopoly* model was studied in Cellini and Lambertini [2009]. Assuming symmetric firms, and identical initial technologies, it follows that there exists a symmetric Nash equilibrium. Therefore, the duopoly analytically reduces to a single-agent optimization since $q_i = q(t)$ and $a_i = a(t)$ for all i , yielding in turn $c_i(t) = c(t)$ equal production costs for all times. Thus, both firms pursue identical production and R&D strategies.

7 Discussion and Extensions

7.1 Comparison to Related Models

We note that the non-negativity constraints $q_i(n) \geq 0, a_i(n) \geq 0$ are crucial in many of the features that emerge. Indeed, they lead to different equilibria which due to the coupling between game stages affect the global strategies of the players. In contrast, in the central planner context of Lafforgue [2008], the same constraints were never binding so that the single agent was always producing, consuming and innovating for all times. (In fact the analytic solution obtained in Lafforgue [2008] is fully stationary, so that the effort allocation is constant in t ; in contrast our controls are fundamentally dependent on current stage $\mathbf{N}(t)$.)

Moreover, the system (22) is similar to the discretized central planner problem of Lafforgue [2008], except that in the latter paper there is no explicit price or production cost but rather a production function entering a consumption social benefit (utility) function. Thus, the analogue of the decreasing production cost $c_1(n)$ is taken up by the productivity factor $B(t)$ which grows proportionally with each innovation. By selecting a Cobb-Douglas production and power utility, Lafforgue [2008] is able to construct an explicit analytic solution to (22).

7.2 Effect of Competition

Rivalry considerations are central to our game framework. In comparison to the single-agent monopoly model, competition lowers both profits and demand elasticity. Indeed the monopolist can capture the full consumer surplus and is therefore highly sensitive to her production costs, encouraging R&D. In contrast, so far we assumed perfectly substitutable goods so that in a duopoly profits are shared in proportion to the current production costs. To interpolate between these two extremes, we next consider the intermediate case of a differentiated Cournot duopoly with imperfect substitution (which goes back at least to [Singh and Vives, 1984]; see also [Ledvina and Sircar, 2012]). Let $0 \leq \gamma \leq 1$ be the substitutability fraction. We assume a linear price function with saturation point $\eta = 1$, so that the price obtained by player i is

$$P_i(q_i, Q_{-i}) = 1 - q_i - \gamma Q_{-i}. \quad (28)$$

Thus, $\gamma = 1$ is perfect substitutability, $P(Q) = 1 - Q$, and $\gamma = 0$ means perfectly complementary goods, $P_i(q_i) = 1 - q_i$, so that within their respective markets players act independently as de facto monopolies.

We first briefly summarize the static equilibrium ignoring R&D. As before, agents choose q_i to maximize $\pi_i = q_i \cdot (P_i(q_i, Q_{-i}) - c_i)$ which leads to the static solution

$$q_i^{Stat} = \frac{(2 - \gamma) + \gamma c_j - 2c_i}{(4 - \gamma^2)} \quad (29)$$

and revenue $\pi_i(\mathbf{c}) = (q_i(\mathbf{c}))^2$. The R&D equilibrium is achieved exactly as before, solving the same system (17) but with the latter modified profits $\pi_i(\mathbf{c})$.

Figure 8 presents the numerical solution at a fixed stage (\cdot, n_2) using the technology ladder $c(n) = e^{-0.05n}$ and a range of γ 's. As expected, lower competition (i.e. lower γ) increases game values and raises production. However, counterintuitively the impact on R&D is ambiguous. Depending on n_2 , less competition may either increase or decrease $a_1^*(n)$. Moreover, the time profile of R&D shifts. Less competition induces higher R&D expenditures up-front (when the near-monopolist is more sensitive to reduced costs), but lower R&D expenditures in the middle and end of progress (when the duopoly affords additional gains at the expense of player 2). Thus, overall the expected time to go from say stage 1 to stage n_1 need not be monotone in γ . In Figure 8 we see that in an undifferentiated market ($\gamma = 1$), a producer might get completely discouraged when far behind $c_1 \gg c_2$ and give up on R&D altogether: $a_1^*(n_1) = 0$ for n_1 small enough. In contrast, with differentiated goods, the producer is more sensitive to her costs and will engage in R&D even when $c_1 \simeq 1$.

7.3 R&D Effort Complementary to Production

In economic growth theory one views labor as the single source of productive capacity which must then be allocated to different activities. From that perspective, choices about R&D investment and production level are not independent, being both labor-consuming. The resulting link couples the optimization problem arising in (17) across the $a(t)$ and $q(t)$ controls. Under strong coupling this




Figure 8: Comparison of R&D effort levels $a_1^*(\cdot, n_2)$ in a dynamic bilateral R&D game with differentiated goods. We take an exponential technology ladder $c_i(n) = \exp(-0.05n)$, $n = 1, \dots, 50$ with $\bar{\lambda} = 10$, $r = 0.1$ and $\mathcal{R}(a) = a^2/2 + 0.1a$ (players are fully symmetric). The price function is linear as in (28) and results are shown for $\gamma \in \{0.4, 0.8, 1\}$. The latter case is the perfectly substitutable goods. Throughout we use $n_2 = 10$ (i.e. $c_2(n_2) = 0.61$).

link is direct, $G(a(t), q(t)) = \text{const}$, see e.g., Lafforgue [2008]. Under a weak coupling $G(a(t), q(t)) \leq \text{const}$ there is only a constraint which may or may not be binding (e.g. labor could be saved, invested in other sectors, etc.).

To illustrate the impact of coupling q and a we briefly revisit the linear duopoly case with unilateral R&D by P1 where we assume that the set of admissible strategies is $\mathcal{A}_i(\mathbf{n}) = \{(q, a) : q + a = A\}$. Thus, R&D and production are perfect substitutes, and investing in R&D entails lower production. This generates an intrinsic shadow cost to R&D and hence we take $\mathcal{R}(a) \equiv 0$. In this scenario, the optimization problems are given by:

$$\begin{aligned} v_1(n; q_1) &= \sup_{q_1(t)} \mathbb{E}_n \left\{ \int_0^\infty e^{-\rho_1 t} q_1(t) [1 - q_1(t) - q_2^*(t) - c_1(N(t))] dt \right\} \\ &= \sup_{q_1 \in [0, A]} \frac{\bar{\lambda}[A - q_1]}{\bar{\lambda}[A - q_1] + \rho_1} [(q_1(1 - q_1 - q_2^*(n) - c_1(n)) + \bar{\lambda}(A - q)v_1(n + 1)], \end{aligned}$$

where (N_t) with $N_0 = 0$ is the counting process for the number of technology advances of P1 by time t .

Either of the double constraints $0 \leq q_1(n) \leq A$, may be binding. If $q_1(n) = 0$ then all labor is devoted to R&D while production is shut-down; if $q_1(n) = A$ then all labor is devoted to production with no R&D. The complementarity between production and R&D introduces additional strategic effects into the dynamic game. Clearly, R&D becomes more “expensive”, and consequently may be shut-down sooner (or started later), making more game stages absorbing. Furthermore, a player may strategically lower production to devote more effort to R&D. As a result, the fixed market structure transitions (which were fully determined by $\mathbf{c}(\mathbf{n})$) become endogenous. First, it is possible that $q_1(\mathbf{n}) = 0$ even if $q_1^{Stat}(\mathbf{n}) > 0$. In other words, there may be a strategic delay by P1 in entering the market, even if the present costs support a duopoly, in order to devote resources exclusively to R&D. Second, it is possible in the later stages that $q_2(\mathbf{n}) > 0$ even if $q_2^{Stat}(\mathbf{n}) = 0$, so that player 2 is allowed to remain in the market even when P1 has costs low enough to become a monopolist.

This happens because monopolist is very sensitive to her costs and therefore could be willing to delay her monopoly transition by tolerating “zombie” competition in order to lower production costs even quicker.

7.4 Spillovers

It may be viewed as unrealistic to assume that the two players compete in a single market (i.e. the goods produced are perfectly substitutable), yet their technology states are completely independent, so that technology gains are fully internalized by the innovator. One could patch this concern by introducing spill-over effects. Spillovers are well-documented empirically and tend to lower R&D investments and therefore reduce productivity growth. Game-theoretically, spillovers can be viewed as either raising the innovation rate of competitors in a static set-up, or removing first-mover advantages after innovation success, see for example [Cellini and Lambertini, 2009, Dawid et al., 2013].

In our model, one can incorporate both immediate and gradual spill-overs. First, one could allow for a possibility of joint jumps in the different components of \mathbf{N} , i.e. common technology innovation based on total R&D investment. This would introduce terms of the form $v_1(n_1 + 1, n_2 + 1)$ into the system (17). Second, one could allow the R&D hazard rates λ_i to be a general function of $\mathbf{a}(t)$, so that R&D investment of player i not only speeds up innovation of player i but also of other players j . Third, one could make $\lambda_i(t)$ to depend on $\mathbf{N}(t)$ to model long-term spill-overs: e.g. if $N_2(t) \gg N_1(t)$, one could imagine that catching up becomes easier for P1 so that $\lambda_1(t)$ increases. Combining these features one could model duopoly R&D innovation rates as e.g.,

$$\lambda_1(a_1, a_2, N_1, N_2) = b_1 a_1 + b_2 a_2 + b_3 (N_2 - N_1)_+,$$

for some constants $b_1, b_2, b_3 \in \mathbb{R}_+$. Yet another possibility is to couple the players’ R&D costs (through either crowding out of R&D labor/resources or market synergies) by making $\mathcal{R}_i(a_i)$ depend on the full profile of efforts \mathbf{a} . Incorporating all these extensions into (17) is straightforward, requiring just an appropriate modification to the local optimization problems.

8 Conclusion

We have presented a framework for modeling oligopolistic stochastic R&D games. In our set-up R&D successes are discrete, allowing to explicitly describe the local Markov Nash equilibria and providing a direct link to the classical static Cournot games. As a result, the market state is represented as a Markov chain coupling the producers. We show that constructing a dynamic equilibrium only requires numerically solving a sequence of nonlinear equations which can be done efficiently in a rather general setting.

Our findings show the multiple nuances of endogenizing market structure. First, we demonstrate that a variety of markets can coexist in a single model (e.g. investing in R&D but not producing, producing but not investing, etc.). Second, we find that the phase transitions induced by structural changes (such as transition to monopoly) generate new dynamic effects, such as non-monotone R&D trends. Third, we find that the impact of uncertainty is double-edged, lowering flexibility but also increasing R&D NPV. Thus, the effect of lowering R&D variance is ambiguous both locally (in terms of immediate R&D effort) and globally (making some states absorbing or not). Fourth, we find that competition (measured either through the production costs of competitors or degree of substitutability) also has ambiguous effects on R&D.

The presented model was motivated by the R&D history in the oil (or more generally energy) market. In that context, a key idea is that of exhaustibility— the running out of *reserves* because the underlying resource is non-replenishable. Exhaustibility adds shadow costs to extraction and introduces a different mechanism for dynamic market structure— see Harris et al. [2010], Ludkovski

and Sircar [2011], Ledvina and Sircar [2012]. Combining exhaustibility and R&D (which would provide insights into the empirical cycles of energy exploration and production) requires introduction of an additional state variable (X_t) for reserves and will be dealt with in a separate work.

Another fruitful research direction is a further study regarding the impact of the number of players L competing. To handle different market sizes in a unified matter, one can represent the current market structure in terms of a probability measure μ_t which keeps track of the cost profile \mathbf{c} via $\mu_t = \frac{1}{L} \sum_{i=1}^L \delta_{c_i(t)}$. One can then index all other game quantities, such as the profits $\pi_i(\mu_t)$ via μ_t . The limiting case $L \rightarrow \infty$ would connect to the burgeoning literature on mean-field games. See, for instance, Chan and Sircar [2015] for some results on Cournot mean field games of exhaustible producers.

Appendices

A Proof of Lemma 1

Proof. If player i is not part of the equilibrium, $q_i^* = 0$, then clearly she's insensitive to either her costs or her competitors' costs. Similarly, if $q_j^* = 0$ then player j is sitting out and hence her costs c_j cannot affect player i , $\frac{\partial q_i^*}{\partial c_j} = 0$. It follows that we can restrict attention to the active set of ℓ^* players and assume that the corresponding equilibrium is of interior type, $q_i^* > 0 \forall i \leq \ell^*$ (the case where a player is just on the boundary of the active set can be treated similarly). With that assumption, differentiating the first-order-conditions defining q_i^* in (6) with respect to c_j we obtain:

$$0 = \frac{\partial q_i}{\partial c_j} P'(Q^*) + [q_i P''(Q^*) + P'(Q^*)] \cdot \frac{\partial Q}{\partial c_j} - 1_{\{i=j\}}. \quad (30)$$

Adding up the expressions in (6) across $i = 1, 2, \dots, \ell$ we obtain

$$\begin{aligned} 0 &= \sum_{i=1}^{\ell} \left(\frac{\partial q_i}{\partial c_j} P'(Q^*) + [q_i P''(Q^*) + P'(Q^*)] \cdot \frac{\partial Q}{\partial c_j} - 1_{\{i=j\}} \right) \\ &= \frac{\partial Q}{\partial c_j} \cdot [P'(Q) + Q P''(Q) + \ell P'(Q)] - 1. \end{aligned} \quad (31)$$

Using the fact that $Q P''(Q) = -\rho P'(Q)$ the last equation simplifies to

$$P'(Q) \frac{\partial Q}{\partial c_j} = \frac{1}{[\ell^* + 1 - \rho]}, \quad (32)$$

which is independent of c_j . In particular we see that total production decreases with respect to any c_j as long as $\rho < \ell^* + 1$. To make sure that this is true even for a monopolist, we require $\rho < 2$. Returning to equation (30) and substituting $P''(Q) = -\rho P'(Q)/Q$ as well as (32) we obtain

$$\begin{aligned} \frac{\partial q_i}{\partial c_j} P'(Q) - \rho \frac{q_i}{Q} P'(Q) \frac{\partial Q}{\partial c_j} + P'(Q) \frac{\partial Q}{\partial c_j} &= 1_{\{i=j\}} \\ \frac{\partial q_i}{\partial c_j} P'(Q) &= 1_{\{i=j\}} - \frac{1 - \rho \frac{q_i}{Q}}{\ell^* + 1 - \rho}. \end{aligned}$$

Solving the last equation for $\frac{\partial q_i}{\partial c_j}$ in each of the cases $i = j$ and $i \neq j$ yields (12). Note again that the impact of any other production costs c_j on q_i is independent of j due to the symmetry imposed by aggregate production. \square

B Proof of Lemma 2

Proof. For concreteness, we consider P1 first. Let $c_1 \geq 0$ be the largest accumulation point of the ladder \mathcal{C}_1 . Using the boundedness and monotonicity of the revenues in terms of costs, we have that for any $\varepsilon > 0$,

$$|\pi_1^*(c_1, c_2(n_2)) - \pi_1^*(c_1(n_1), c_2)| \leq \varepsilon \rho$$

for all (n_1, n_2) large enough. Therefore, since

$$\frac{\pi_1^*(c_1, c_2(n_2))}{\rho_1} \geq v_1(n_1, n_2) \geq \frac{\pi_1^*(c_1(n_1), c_2)}{\rho_1} \quad (33)$$

(which holds since starting at (n_1, n_2) , the left- and right-most numerators express the minimum and maximum revenue rates possible going forward) we have

$$|\pi_1^*(c_1(n_1), c_2(n_2)) - \rho_1 v_1(n_1 + 1, n_2)| \leq \varepsilon \quad \text{and} \quad |v_1(n_1, n_2 + 1) - v_1(n_1 + 1, n_2)| \leq \varepsilon \quad (34)$$

for all $n_1 \geq N_1(\varepsilon), n_2 \geq N_2(\varepsilon)$ as above.

Now fix $n_1 \geq N_1, n_2 \geq N_2$ and the corresponding $a_2^*(\mathbf{n})$. As notational shorthand we write $\pi \equiv \pi_1^*(\mathbf{c}(\mathbf{n}))$. We wish to compare the expected reward to P1 from investing in R&D at some level $a_1 > 0$ at the current stage, i.e. $J_1(a_1)$ from (20), vis-a-vis zero R&D investment, $J_1(0)$. In particular, we want to show that the latter is advantageous, i.e. $J_1(0) - J_1(a_1) \geq 0$. Plugging in $a_1 = 0$ into (20) yields $J_1(0) - J_1(a_1) =$

$$\begin{aligned} & \frac{1}{a_2^* + \rho_1} [\pi + a_2^* v_1(n_1, n_2 + 1)] - \frac{1}{a_1 + a_2^* + \rho_1} \{ \pi - \rho_1 \mathcal{R}_1(a_1) + a_1 v_1(n_1 + 1, n_2) + a_2^* v_1(n_1, n_2 + 1) \} \\ &= \{ \pi + a_2^* v_1(n_1, n_2 + 1) \} \left\{ \frac{1}{a_2^* + \rho_1} - \frac{1}{a_1 + a_2^* + \rho_1} \right\} + \frac{\rho_1}{a_1 + a_2^* + \rho_1} \mathcal{R}_1(a_1) - \frac{a_1}{a_1 + a_2^* + \rho_1} v_1(n_1 + 1, n_2) \\ &\geq \{ \pi + a_2^* v_1(n_1, n_2 + 1) \} \frac{a_1}{(a_1 + a_2^* + \rho_1)(a_2^* + \rho_1)} + \frac{\rho_1 a_1}{a_1 + a_2^* + \rho_1} \mathcal{R}'_1(0) - \frac{a_1}{a_1 + a_2^* + \rho_1} v_1(n_1 + 1, n_2) \\ &= \frac{a_1}{a_1 + a_2^* + \rho_1} \left\{ \frac{\pi - \rho_1 v_1(n_1 + 1, n_2)}{a_2^* + \rho_1} + \rho_1 \mathcal{R}'_1(0) + \frac{a_2^*(v_1(n_1, n_2 + 1) - v_1(n_1 + 1, n_2))}{a_2^* + \rho_1} \right\}, \quad (35) \end{aligned}$$

where the inequality is based on the convexity of $\mathcal{R}_1(a)$, $\mathcal{R}_1(a_1) \geq a_1 \mathcal{R}'_1(0)$. Now both the first and last terms (that measure respectively the gain from a R&D success and a loss from R&D failure) are by assumption less than ε , while the middle term $\rho_1 \mathcal{R}'_1(0)$ is by assumption strictly positive. We conclude

$$J_1(0) - J_1(a_1) \geq \frac{a_1}{a_1 + a_2^* + \rho_1} \left\{ -\frac{\rho_1 \varepsilon}{a_2^* + \rho_1} + \rho_1 \mathcal{R}'_1(0) - \frac{a_2^* \varepsilon}{a_2^* + \rho_1} \right\} = \frac{a_1}{a_1 + a_2^* + \rho_1} \{ -\varepsilon + \rho_1 \mathcal{R}'_1(0) \} \geq 0$$

as soon as $\varepsilon < \rho_1 \mathcal{R}'_1(0)$. The case for other players is dealt with similarly. \square

C Proof of Lemma 3

Proof. Intuitively, the refined model offers higher value since it gives the players a more fine control over their R&D efforts. Indeed, one can always “embed” a model with M stages into one with $2M$ stages by keeping efforts constant for pairs of consecutive technology levels, $\check{a}^{(2M)}(2k) = \check{a}^{(2M)}(2k + 1) = a^{(M)}(k)$. Optimization over a larger control set should increase profits.

However, refinement also removes the opportunity for rapid advancement, which lowers the NPV of future profits. To wit, consider the first innovation date $\tau \equiv \tau^{(M),1}$ in the model with M stages with an effort level $a \equiv a_1(1)$. Then, $\tau \sim \text{Exp}(M\bar{\lambda}a)$ and the corresponding discount

factor is $\mathbb{E}[e^{-\rho_1\tau}] = \frac{M\bar{\lambda}a}{M\bar{\lambda}a + \rho_1}$. In the refined model with $2M$ stages, the comparable epoch is the *second* innovation date $\tau^{(2M),2}$. Suppose that same effort is maintained $\check{a}^{(2M)}(1) = \check{a}^{(2M)}(2) = a$ on $[0, \tau^{(2M),2})$. Then times to success on each sub-stage, $\tau^{(2M),1}, \tau^{(2M),2} - \tau^{(2M),1} \sim \text{Exp}(2M\bar{\lambda}a)$ are i.i.d., so that $\tau^{(2M),2} \sim \text{Ga}(2, 2M\bar{\lambda}a)$ has a Gamma distribution. Therefore, the discount factors satisfy

$$\mathbb{E}\left[e^{-\rho_1\tau^{(2M),2}}\right] = \left(\frac{2M\bar{\lambda}a}{2M\bar{\lambda}a + \rho_1}\right)^2 < \frac{M\bar{\lambda}a}{M\bar{\lambda}a + \rho_1} = \mathbb{E}[e^{-\rho_1\tau^{(M),1}}] \quad \forall a. \quad (36)$$

Observe that $\tau^{(2M),2}$ is a mean-preserving spread of $\tau^{(M),1}$ with a smaller variance ($\mathbb{E}[\tau^{(2M),2}] = \mathbb{E}[\tau^{(M),1}] = (M\bar{\lambda}a)^{-1}$, while $\text{Var}(\tau^{(2M),2}) = \frac{1}{2(M\bar{\lambda}a)^2} = \frac{1}{2}\text{Var}(\tau^{(M),1})$), so that (36) follows immediately due to convexity of $t \mapsto \exp(-\rho_1 t)$. Note that the above argument holds for any degree of refinement (not just doubling), since we have that

$$k \mapsto \mathbb{E}\left[e^{-\rho_1\tau^{(kM),k}}\right] = \left(\frac{k\bar{\lambda}a}{k\bar{\lambda}a + \rho_1}\right)^k, \quad \tau^{(kM),k} \sim \text{Ga}(k, k\bar{\lambda}a)$$

is *decreasing* in k .

To formalize the above ideas, consider a setting with unilateral R&D by P1 and a single-stage original model which is refined to a two-stage one. Label the respective profits as $\pi_1(1) < \pi_1(2) < \pi_1(3)$, and respective R&D success times as $\tau^{(1),1}, \tau^{(2),1}, \tau^{(2),2}$. Then similar to (23), fixing the effort levels $a^{(1),1}$ and $a^{(2),1}, a^{(2),2}$ and taking without loss of generality $M\bar{\lambda} \equiv 1$, we obtain that the corresponding expected profits are

$$v_1^{(1)}(1; a^{(1),1}) = \frac{1}{a^{(1),1} + \rho_1} \{\pi_1(1) - \mathcal{R}(a^{(1),1})\} + \frac{a^{(1),1}}{a^{(1),1} + \rho_1} \frac{1}{\rho_1} \pi_1(3); \quad (37)$$

$$\begin{aligned} v_1^{(2)}(1; a^{(2),1}, a^{(2),2}) &= \frac{1}{2a^{(2),1} + \rho_1} \{\pi_1(1) - \mathcal{R}(a^{(2),1})\} + \frac{2a^{(1),1}}{2a^{(1),1} + \rho_1} \frac{1}{2a^{(2),1} + \rho_1} \{\pi_1(2) - \mathcal{R}(a^{(2),2})\} + \\ &+ \frac{2a^{(1),1}}{2a^{(1),1} + \rho_1} \frac{2a^{(2),1}}{2a^{(2),1} + \rho_1} \frac{1}{\rho_1} \pi_1(3). \end{aligned} \quad (38)$$

We now show that depending on the relationship between stagewise profits $\pi_1(n)$, $n = 1, 2, 3$, the ordering between $v_1^{(1)}(1)$ and $v_1^{(2)}(1)$ can turn out either way. This confirms algebraically that the competing effects described above lead to ambiguous outcomes.

$v_1^{(1)}(1) > v_1^{(2)}(1)$: if the intermediate profits $\pi_1(2)$ are close to $\pi_1(1)$ then there is minimal gain from reduced uncertainty, but less likelihood of a quick breakthrough. So less uncertainty lowers value. Fix the optimal two-stage effort levels $a^{(2),1}, a^{(2),2}$. We construct a dominant R&D strategy for the original model. To do so, take $\tilde{a}^{(1),1}$ such that

$$\frac{\tilde{a}^{(1),1}}{\tilde{a}^{(1),1} + \rho_1} := \frac{2a^{(1),1}}{2a^{(1),1} + \rho_1} \frac{2a^{(2),1}}{2a^{(2),1} + \rho_1}. \quad (39)$$

This is equivalent to picking $\tilde{a}^{(1),1}$ such that $\mathbb{E}[e^{-\rho_1\tau^{(1),1}}] = \mathbb{E}[e^{-\rho_1\tau^{(2),2}}]$. Write $\pi_1(2) = \pi_1(1) + \varepsilon$. Then an algebraic derivation based on plugging in (39) shows that

$$\frac{1}{2a^{(2),1} + \rho_1} \pi_1(1) + \frac{2a^{(1),1}}{2a^{(1),1} + \rho_1} \frac{1}{2a^{(2),1} + \rho_1} \{\pi_1(1) + \varepsilon\} = \frac{1}{\tilde{a}^{(1),1} + \rho_1} \pi_1(1) + C\varepsilon,$$

for a positive constant C . Moreover, a direct computation reveals that in this case NPVs of effort costs satisfy

$$-\frac{1}{2a^{(2),1} + \rho_1} \mathcal{R}(a^{(2),1}) - \frac{2a^{(1),1}}{2a^{(1),1} + \rho_1} \frac{1}{2a^{(2),1} + \rho_1} \mathcal{R}(a^{(2),2}) < -\frac{1}{\tilde{a}^{(1),1} + \rho_1} \mathcal{R}(a^{(1),1})$$

for a quadratic R&D cost function $\mathcal{R}(a)$. Subtracting (37) from (38), it follows that for ε small enough, the two-stage strategy gives strictly lower expected profits, gaining order- ε extra profits from the intermediate sub-stage but losing a strictly positive amount relative to $\tilde{a}^{(1),1}$.

$v_1^{(1)}(1) < v_1^{(2)}(1)$: If the intermediate profits are nearly as large as the terminal ones, $\pi_1(2) = \pi_1(3) - \varepsilon$, then the refined model would lead to higher NPV as it accelerates the appearance of higher profits. To see that, let $\tilde{a} = a^{(1),1}$ be optimal for $v_1^{(1)}(1)$ and set $\tilde{a}^{(2),1} = \tilde{a}^{(2),2} = \tilde{a}$. Substituting into (37)-(38), the resulting difference in P1 expected profits is

$$\begin{aligned} v_1^{(2)}(1) - v_1^{(1)}(1) &= \frac{1}{2\tilde{a} + \rho_1} \{ \pi_1(1) - \mathcal{R}(\tilde{a}) \} + \frac{2\tilde{a}}{(2\tilde{a} + \rho_1)^2} \{ \pi_1(2) - \mathcal{R}(\tilde{a}) \} + \frac{2\tilde{a}^2}{(2\tilde{a} + \rho_1)^2} \frac{1}{\rho_1} \pi_1(3) \\ &\quad - \frac{1}{\tilde{a} + \rho_1} \{ \pi_1(1) - \mathcal{R}(\tilde{a}) \} - \frac{\tilde{a}}{\tilde{a} + \rho_1} \frac{1}{\rho_1} \pi_1(3) \\ &= -\tilde{a}\rho_1 \mathcal{R}(\tilde{a}) + (2\tilde{a} + \rho_1) \cdot \tilde{a} \cdot (\pi_1(3) - \pi_1(1)) - C\varepsilon. \end{aligned}$$

Hence, for $\pi_1(3) - \pi_1(1)$ large enough and ε small enough, we have that the 2-stage NPV is larger. \square

Figure 6 confirms that both of the above outcomes arise. In the Figure, for $c_1(x)$ large or very small, we are in the second situation of $\pi_1(2) \simeq \pi_1(1)$ (since profits are currently small, or there are little productivity gains left) and $v_1^{(M)}(x) > v_1^{(2M)}(x)$. For $c_1(x)$ moderate, there is more sensitivity to costs and $v_1^{(M)}(x) < v_1^{(2M)}(x)$ (just barely in the plot).

References

- A. Azevedo and D. Paxson. Developing real option game models. *European Journal of Operational Research*, 237(3):909–920, 2014.
- P. Bremaud. *Point Processes and Queues*. Springer, New York, 1981.
- R. Cellini and L. Lambertini. Dynamic R&D with spillovers: Competition vs cooperation. *Journal of Economic Dynamics and Control*, 33(3):568–582, 2009.
- P. Chan and R. Sircar. Bertrand & Cournot mean field games. *Applied Mathematics & Optimization*, 71(3):533–569, 2014.
- P. Chan and R. Sircar. Fracking, renewables & mean field games. Technical report, Princeton University, 2015. Available at SSRN 2632504.
- E. Dahlgren and T. Leung. An optimal multiple stopping approach to infrastructure investment decisions. *Journal of Economic Dynamics and Control*, 53:251–267, 2015.
- M. H. A. Davis. *Markov Models and Optimization*. Chapman & Hall, London, 1993.
- H. Dawid, M. Kopel, and PM Kort. R&D competition versus R&D cooperation in oligopolistic markets with evolving structure. *International Journal of Industrial Organization*, 31(5):527–537, 2013.

- U. Doraszelski. An R&D race with knowledge accumulation. *The RAND Journal of Economics*, pages 20–42, 2003.
- G. Femminis and G. Martini. Irreversible investment and R&D spillovers in a dynamic duopoly. *Journal of Economic Dynamics and Control*, 35(7):1061–1090, 2011.
- S. Fölster and G. Trofimov. Industry evolution and R&D externalities. *Journal of Economic Dynamics and Control*, 21(10):1727–1746, 1997.
- D. Fudenberg and J. Tirole. Preemption and rent equalization in the adoption of new technology. *The Review of Economic Studies*, 52(3):383–401, 1985.
- D. Fudenberg, R. Gilbert, J. Stiglitz, and J. Tirole. Preemption, leapfrogging and competition in patent races. *European Economic Review*, 22(1):3–31, 1983.
- L.H. Goulder and S.H. Schneider. Induced technological change and the attractiveness of CO₂ abatement policies. *Resource and Energy Economics*, 21(3):211–253, 1999.
- M. Goyal and S. Netessine. Strategic technology choice and capacity investment under demand uncertainty. *Management Science*, 53(2):192–207, 2007.
- A. Grimaud, G. Lafforgue, and B. Magné. Climate change mitigation options and directed technical change: A decentralized equilibrium analysis. *Resource and Energy Economics*, 33(4):938–962, 2011.
- G.M. Grossman and C. Shapiro. Dynamic R&D competition. *The Economic Journal*, pages 372–387, 1987.
- O. Guéant, J.-M. Lasry, and P.-L. Lions. Mean field games and applications. In *Paris-Princeton Lectures on Mathematical Finance 2010*, pages 205–266. Springer, 2011.
- C. Harris and J. Vickers. Racing with uncertainty. *The Review of Economic Studies*, 54(1):1–21, 1987.
- C. Harris, S. Howison, and R. Sircar. Games with exhaustible resources. *SIAM Journal of Applied Mathematics*, 70:2556–2581, 2010.
- K.J.M. Huisman and P.M. Kort. Strategic technology adoption taking into account future technological improvements: A real options approach. *European Journal of Operational Research*, 159(3):705–728, 2004.
- K.L. Judd. Closed-loop equilibrium in a multi-stage innovation race. *Economic Theory*, 21(2-3): 673–695, 2003.
- M.I. Kamien and N.L. Schwartz. Optimal exhaustible resource depletion with endogenous technical change. *Review of Economic Studies*, 45(1), 1978.
- C. Krauss. Drillers answer low oil prices with cost-saving innovations. *New York Times*, (May 12), 2015. URL <http://www.nytimes.com/2015/05/12/business/energy-environment/drillers-answer-low-oil-prices-with-cost-saving-innovations.html>.
- G. Lafforgue. Stochastic technical change, non-renewable resource and optimal sustainable growth. *Resource and Energy Economics*, 30(4):540–554, 2008.

- A. Ledvina and R. Sircar. Oligopoly games under asymmetric costs and an application to energy production. *Mathematics and Financial Economics*, 6(4):261–293, 2012.
- M. Ludkovski and R. Sircar. Exploration and exhaustibility in dynamic Cournot games. *European Journal of Applied Mathematics*, 23(3):343–372, 2011.
- R. Pindyck. Uncertainty and exhaustible resource markets. *Journal of Political Economy*, 88 (6): 1203–1225, 1980.
- J.F Reinganum. Uncertain innovation and the persistence of monopoly. *The American Economic Review*, pages 741–748, 1983.
- J.F Reinganum. The timing of innovation: Research, development, and diffusion. In *Handbook of industrial organization*, pages 849–908. Elsevier, 1989.
- N. Singh and X. Vives. Price and quantity competition in a differentiated duopoly. *The RAND Journal of Economics*, pages 546–554, 1984.
- Y. Tsur and A. Zemel. Optimal transition to backstop substitutes for nonrenewable resources. *Journal of Economic Dynamics and Control*, 27(4):551–572, 2003.
- X. Vives. *Oligopoly pricing: old ideas and new tools*. The MIT press, 2001.
- H. Weeds. Strategic delay in a real options model of R&D competition. *Review of Economic Studies*, 69(3):729–747, 2002.