

# The timing of version releases: A dynamic duopoly model

Ron N. Borkovsky<sup>1</sup>

Received: 21 October 2015 / Accepted: 26 June 2017 / Published online: 19 August 2017  
© Springer Science+Business Media, LLC 2017

**Abstract** In many R&D-intensive consumer product categories, firms deliver value to consumers through the quality enhancements provided by new and improved versions of existing products. Therefore, important marketing decisions relate to a firm's strategy for developing quality enhancements and releasing new versions. This paper explores this type of product development using a dynamic duopoly model that endogenizes each firm's decisions over how much to invest in R&D and when to release new versions. Specifically, I explore how two key industry fundamentals—the degree of horizontal differentiation and the cost of releasing a new version—affect firms' product development strategies and, accordingly, the evolution of industry structure. I find that varying the degree of horizontal differentiation gives rise to three distinctly different types of competitive dynamics: *preemption races* when the degree of horizontal differentiation is low; *phases of accommodation* when it is moderate; and *asymmetric R&D wars* when it is high. Furthermore, I find that an increase in the cost of releasing a new version can induce firms to compete *more* aggressively for the lead and, in doing so, release new versions *more* frequently despite the higher cost.

**Keywords** Product development · Upgrades · Research and development · Innovation · Quality ladder model · Dynamic oligopoly

**JEL Classification** C73 · D43 · L13 · M31

---

✉ Ron N. Borkovsky  
ron.borkovsky@rotman.utoronto.ca

<sup>1</sup> Rotman School of Management, University of Toronto, Toronto, ON, Canada

## 1 Introduction

In many R&D-intensive consumer product categories, firms invest heavily in R&D in order to develop improvements for their existing products. They bring these improvements to market by periodically releasing new versions of their products that incorporate them. This is true of consumer electronics (e.g., smart phones, tablet computers, and video game consoles), software (e.g., office software suites and web browsers), and R&D-intensive consumer-packaged goods (CPG) categories, among others.

Consider the diapers category. Since the 1970s, Procter & Gamble (Pampers, Luvs) and Kimberly-Clark (Huggies) have intensely waged the so-called diaper war by investing heavily in R&D to develop improvements in comfort, absorbency, and containment, and bringing these improvements to market via version releases (Elzinga and Mills 1996; Parry and Jones 2001; Dyer et al. 2004). Both P&G and Kimberly-Clark have developed and incorporated reusable tabs, elastic leg bands, gel technology, and breathable material, among other improvements. This competition to innovate continues even today, as P&G and Kimberly-Clark continue to invest heavily in R&D to improve fit, absorbency, durability, and odor protection (Alfonsi et al. 2010). Other R&D-intensive CPG categories—such as toothpaste<sup>1</sup>—have exhibited similar competitive dynamics.

In such categories, firms deliver value to consumers through the quality enhancements (or new “features”) provided by new versions. Therefore, important marketing decisions relate to a firm’s strategy for the development of product improvements and the timing of version releases. While they enhance demand, version releases are costly. Hence a firm may choose to accumulate new innovations over time and release a new version that incorporates them only periodically. Such categories are often concentrated because the high costs associated with both the R&D process and advertising serve as barriers to entry (Sutton 1991). Therefore, it is important to account for the strategic interaction that inevitably characterizes firms’ R&D investment and version release strategies. However, despite the prevalence of R&D-intensive consumer product categories, the academic literature has thus far given little attention to oligopolistic competition characterized by this type of product development.

In this paper, I seek to better understand how firms decide how much to invest in R&D to develop product improvements and when to release new versions of their products that incorporate them. Moreover, I explore how two industry fundamentals—the degree of horizontal differentiation and the cost of releasing a new version—affect firms’ R&D investment and version release strategies and,

---

<sup>1</sup>P&G (Crest) and Colgate-Palmolive (Colgate) have invested heavily in developing innovations relating to the active ingredient (which fights tooth decay), other therapeutic benefits (e.g., tartar prevention), and packaging, and have introduced these innovations via periodic version releases (McCoy 2001; Parry 2001; Dyer et al. 2004).

accordingly, the evolution of industry structure. I focus on these two industry fundamentals because they play key roles in determining a firm's ability to profit from product development. The degree of horizontal differentiation determines the extent to which firms can translate successful product development into increased share, higher prices, and accordingly higher profits. The cost of releasing a new version determines the extent to which the cost of bringing product improvements to market erodes the increased profits that these product improvements can generate.

I model product development within the context of the Ericson and Pakes (1995) framework for numerically analyzing dynamic models of oligopolistic competition—in particular, dynamic stochastic games. (See Doraszelski and Pakes (2007) for a detailed review.) My model differs from earlier quality ladder oligopoly models (e.g., Pakes and McGuire 1994; Borkovsky et al. 2012) by allowing each firm to accumulate R&D successes (or “stockpile R&D”) over time and then decide exactly when to release a new version into which it attempts to incorporate them. Firms also engage in static price competition, where the degree of vertical product differentiation is determined by the firms' (endogenous) product qualities, and the degree of horizontal differentiation is exogenous. Finally, the model restricts attention to the business stealing effect and in doing so abstracts from the market expansion effect in order to focus on how firms use R&D to improve competitive positioning. While the model could be used to study any of the product categories described above, because the model of price competition is static and therefore does not explicitly incorporate product durability, it perhaps best applies to R&D-intensive CPG categories.<sup>2</sup>

The Ericson and Pakes (1995) framework is well-suited for studying the type of product development described above. First, the framework can easily accommodate a model in which firms make numerous successive releases, whereas the analytic theory literature on product launches tends to focus on new product introductions and therefore restricts attention to a single release decision per product.<sup>3,4</sup> Second, to study this type of product development, one must keep track of each firm's product quality and *R&D stock*—the stock of product improvements that it has developed since its last version release. This demands a model with a multi-dimensional state space. However, dynamic stochastic games with multi-dimensional state spaces tend

---

<sup>2</sup>Goettler and Gordon (2011) incorporate product durability into a quality ladder model in the Ericson and Pakes (1995) framework and use it to study competition between Intel and AMD. While their model endogenizes firms' R&D spending decisions, it does not endogenize the timing of version releases; rather, it assumes that a firm releases a new product after each successful innovation.

<sup>3</sup>See Scherer (1967), Kamien and Schwartz (1972), Wilson and Norton (1989), Moorthy and Png (1992), and Bhaskaran and Ramachandran (2011).

<sup>4</sup>Hitsch (2006) devises and estimates a model of a firm's decisions on whether to launch—and subsequently whether to scrap—new products in the face of demand uncertainty that resolves itself over time (post-launch) as sales are realized. The model differs from the one in this paper in that it (i) focuses on new products with unchanging quality and therefore incorporates neither R&D investment nor repeated releases (but does incorporate advertising); and (ii) does not incorporate strategic interaction between firms.

to be analytically intractable.<sup>5,6</sup> The Ericson and Pakes (1995) framework is well-suited for this problem because it allows for numerical analysis of models with multi-dimensional state spaces. Third, there are a few papers that explicitly consider repeated releases of the same product, but they either do so in a monopoly context (Ramachandran and Krishnan 2008) or they exogenize the behavior of rival firms (Morgan et al. 2001; Aizcorbe 2005). By numerically analyzing a model in the Ericson and Pakes (1995) framework, I am able to study such product development in an oligopolistic context. The equilibrium behaviors that I uncover (described in the next paragraph) are inherently oligopolistic and therefore could not possibly arise in a single-agent model. In summary, by numerically analyzing a model in the Ericson and Pakes (1995) framework, I am able to study an important type of product development that is not accessible through more familiar analytic research methods.

I find that varying the degree of horizontal differentiation gives rise to three distinctly different types of competitive dynamics. First, when the degree of horizontal differentiation is low, firms engage in *preemption races*; when neither firm has too large a lead, firms compete aggressively for the lead by investing heavily in R&D and releasing new versions frequently, and one firm eventually comes to dominate the market. Second, when the degree of horizontal differentiation is moderate, firms enter *phases of accommodation* during which they compete less aggressively. Accommodation is possible only because the leader induces the follower to sharply reduce its R&D investment by threatening to release a new version—which would incorporate its superior R&D stock—if the follower does not comply. Accommodation slows the leader's march to market dominance. Third, when the degree of horizontal differentiation is high, firms engage in *asymmetric R&D wars*; i.e., they engage in aggressive competition once one firm achieves a sufficiently large lead. The understanding that aggressive competition would erupt were one firm to gain a large lead induces both firms to avoid striving for a large lead, hence neither firm comes to dominate the market. Finally, I establish that phases of accommodation and asymmetric R&D wars are unique to the setting I study in the sense that they arise *only* if firms can stockpile R&D and decide exactly when to release new versions.

The results summarized above emphasize the strategic nature of R&D stockpiling; in each of the three scenarios, a firm builds an R&D stock not only because it can ultimately enhance profits, but also because it can be used to influence rival behavior. (See Section 6.1 for details.)

---

<sup>5</sup>Aoki (1991), Harris (1991), Budd et al. (1993), and Hörner (2004) explore R&D competition using analytically tractable dynamic stochastic games. To achieve analytic tractability, each assumes that there is a one-dimensional state space that reflects the size of one firm's lead. While I too restrict attention to the size of a firm's lead in the product market, my research question cannot be explored using an analytically tractable model because I require two additional states that track the firms' respective R&D stocks.

<sup>6</sup>Ofek and Sarvary (2003) devise an analytic model (with a one-dimensional state indicating whether a firm is leader or follower) to study dynamic competition in markets in which firms invest in R&D to develop next-generation products. They explore the implications of different advantages that a leader might possess in terms of innovative ability, reputation, and advertising effectiveness. This paper differs from Ofek and Sarvary (2003) in that it focuses on the strategic role of R&D stockpiling and (endogenous) version releases.

I also explore the effect of the cost of releasing a new version (hereafter, the “release cost”) on firms’ product development strategies. Because an increase in the release cost erodes the profits that can be earned from future versions, one might expect it to induce firms to slow the pace of product development. However, I find that the opposite can occur; i.e., an increase in the release cost can induce firms to release new versions *more* frequently despite the higher cost. This perhaps counterintuitive result arises because as the release cost increases, it becomes more costly for a product market follower to catch up to the leader. This weakens the follower’s incentive to innovate and, as a result, the follower is less likely to ever catch up to the leader. So, a higher release cost makes the position of a product market leader more secure, which induces firms to compete *more* aggressively for the lead.

In addition to generating the insights described above, my model and results could be very useful for future empirical work. Dynamic oligopoly models have played a central role in recent empirical work in marketing and industrial organization; see Aguirregabiria and Nevo (2013) and Bajari et al. (2013). It can however be very difficult to adapt a given dynamic oligopoly model to a specific empirical context without first having a deep understanding of the equilibrium behaviors that the model admits. This presents a challenge because one cannot realistically aspire to understand such a model without thoroughly exploring it, and this is especially true if the model includes multiple choice variables—in this case, R&D investment and version release decisions—which can interact in complex ways.

Consider the following two examples of empirical papers that are based on computational applied theory papers that preceded them: Qi (2013) adapts the Doraszelski and Markovich (2007) dynamic model of goodwill advertising to explore the impact of the 1971 U.S. cigarette advertising ban on industry structure, and is able to corroborate their findings. Borkovsky et al. (2017) exploit the in-depth understanding of the quality ladder model that Borkovsky et al. (2012) deliver in order to adapt it to study brand management and measure brand value in the stacked chips category. This paper could play a similar role in facilitating empirical work on R&D-intensive consumer product categories. In Section 7, I discuss several product categories to which the model could be applied, as well as the ways in which one might tailor the model to address the challenges that those product categories present.

The paper proceeds as follows. Sections 2 and 3 describe the model and my computational strategy. Sections 4–6 present the results. Section 7 explains how this paper helps us better understand product development in categories characterized by repeated releases; discusses limitations and possible extensions; and concludes.

## 2 The model

I model product development within the context of a dynamic quality ladder duopoly (Pakes and McGuire 1994; Borkovsky et al. 2012). In this section, I present the model (Section 2.1), discuss two model assumptions (Section 2.2), and describe the equilibrium conditions (Section 2.3).

## 2.1 Model setup

The model is cast in discrete time and has an infinite horizon. A firm engages in product development so as to move up a “quality ladder”. Specifically, in each period, a firm makes two product development decisions. First, it invests in R&D in order to develop product improvements, which it can accumulate over time. Second, it decides whether to release a new version of its product into which it attempts to incorporate the improvements developed since its last version release. Each firm sells a single product; hence, a new version of a product replaces the previous version. Finally, in each period, firms’ engage in static price competition, where the degree of vertical product differentiation is determined by the firms’ respective (endogenous) product qualities, and the degree of horizontal product differentiation is exogenous.

**Firms and states** The state of firm  $i \in \{1, 2\}$  is  $\omega_i \equiv (\omega_i^m, \omega_i^R)$ , where  $\omega_i^m$  is the quality of the good firm  $i$  sells in the product market, and  $\omega_i^R$  is the number of R&D successes that firm  $i$  has achieved—henceforth referred to as its “R&D stock”—since it released the current version of its product. Let  $\omega_i^m \in \Omega^m \equiv \{0, 1, \dots, L^m\}$  and  $\omega_i^R \in \Omega^R \equiv \{0, 1, \dots, L^R\}$ . I devise a model of product market competition (presented below) that allows me to reduce the dimensionality of the state space by restricting attention to the difference between firms’ respective product qualities,  $\omega^m \equiv \omega_1^m - \omega_2^m \in \Omega_d^m \equiv \{-L^m, -L^m + 1, \dots, L^m - 1, L^m\}$ .<sup>7</sup> (I refer to  $\omega^m$  as the product market state.) It follows that the industry state is  $\omega \equiv (\omega^m, \omega_1^R, \omega_2^R) \in \Omega \equiv \Omega_d^m \times \Omega^R \times \Omega^R$ . (I use boldface to distinguish between vectors/arrays and scalars.) Firm  $i$  is able to change its R&D stock over time through investment  $x_i \geq 0$ . Firm  $i$  is able to change the quality of the good it sells in the product market by releasing a new version. When firm  $i$  releases a new version, each unit of its R&D stock is incorporated into the new version with some exogenous probability. This resets firm  $i$ ’s R&D stock,  $\omega_i^R$ , to zero.

In each period, firms first compete in the product market. Each firm then decides whether to release a new version of—or “update”—its product. Finally, firms make R&D investment decisions. Below I elaborate on the sequence of events in each period, describe the static model of product market competition, and then turn to investment and version release dynamics.

**Timing** Each period is divided into two subperiods. Version release decisions occur in subperiod 1 and R&D investment decisions occur in subperiod 2. The sequence of events is as follows.

In subperiod 1:

<sup>7</sup>This approach is common in the theoretical literature on dynamic quality ladder models; see the papers cited in footnote 5.

1. Firms observe the prevailing industry state  $\omega$ . Each firm then learns how much it would cost to release a new version of its product, i.e., it draws a private release cost  $\phi_i \in \Phi$  from a distribution  $G(\cdot)$ .<sup>8</sup>
2. Firms compete in the product market and earn profits.
3. The firms simultaneously decide whether to release new versions.
4. The outcome of the version release process is realized; i.e., if a firm has chosen to release a new version, the release occurs and the firm learns how much of its R&D stock has been successfully incorporated into the new version of its product. The industry state transitions from  $\omega$  to  $\omega'$ , and firms observe the new industry state. If neither firm chooses to release a new version, then  $\omega' = \omega$ .

In subperiod 2:

5. The firms simultaneously decide how much to invest in R&D.
6. The outcomes of investments in R&D are realized. The industry state transitions from  $\omega'$  to  $\omega''$ , and firms observe the new industry state.

**Product market competition** In each period, firms engage in static price competition with products that are differentiated both vertically and horizontally. There is a continuum of consumers. Each consumer purchases exactly one unit of one product, i.e., there is no outside option. As a result, the dynamic model restricts attention to the business stealing effect and abstracts from the market expansion effect. (This assumption is discussed further in Section 2.2.)

The utility a consumer derives from purchasing from firm  $i$  is  $b\omega_i^m - p_i + \sigma\epsilon_i$ , where  $p_i$  is the price,  $\epsilon_i$  represents the consumer’s idiosyncratic preference for product  $i$ , and  $\sigma > 0$ . Assuming that the idiosyncratic preferences  $(\epsilon_1, \epsilon_2)$  are independently and identically distributed as standard type 1 extreme value, the demand for incumbent firm  $i$ ’s product is

$$D_i(\mathbf{p}; \omega_i^m, \omega_{-i}^m) = m \frac{1}{1 + \exp\left(\frac{b(\omega_{-i}^m - \omega_i^m) + p_i - p_{-i}}{\sigma}\right)},$$

where the subscript  $-i$  refers to the rival firm,  $\mathbf{p} = (p_1, p_2)$  is the vector of prices, and  $m > 0$  is the size of the market (the measure of consumers).<sup>9</sup> Like Anderson et al. (1992), I interpret the parameter  $\sigma$  as the degree of horizontal differentiation between products; as  $\sigma$  increases, consumers’ tastes are more heterogeneous, so consumers are less responsive to differences in prices and (vertical) qualities.

<sup>8</sup>I assume that the cost of releasing a new version is private and random because it guarantees the existence of a Markov perfect equilibrium in pure strategies (Doraszelski and Satterthwaite 2010).

<sup>9</sup>In the Appendix, I explain that the Logit demand model can be reinterpreted as an address model (Anderson et al. 1992). I also show that if one replaces the Logit demand model with the Hotelling (1929) address model, the equilibria that the dynamic model admits are qualitatively similar.

The demand that a firm faces is not a function of the firms' absolute product qualities, but only of the difference between them. Therefore, I can write each firm's demand function as a function of  $\omega^m$ :

$$D_1(\mathbf{p}; \omega^m) = m \frac{1}{1 + \exp\left(\frac{-b\omega^m + p_1 - p_2}{\sigma}\right)},$$

$$D_2(\mathbf{p}; \omega^m) = m \frac{1}{1 + \exp\left(\frac{b\omega^m + p_2 - p_1}{\sigma}\right)}.$$

In doing so, I have reduced the dimensionality of the state space.

Incumbent firm  $i$  chooses the price  $p_i$  of its product to maximize profits. Hence, its profits in product market state  $\omega^m$  are

$$\pi_i(\omega^m) = \max_{p_i \in [0, \infty)} D_i(p_i, p_{-i}(\omega^m); \omega^m) (p_i - c),$$

where  $p_{-i}(\omega^m)$  is the price charged by the rival and  $c \geq 0$  is the marginal cost of production. A Nash equilibrium of the product market game (for a given  $\omega^m$ ) is characterized by the system of optimality conditions derived from the firms' respective profit-maximization problems.<sup>10</sup> Because product market competition is static, it does not directly affect state-to-state transitions in the dynamic model; hence, for the purposes of the dynamic model, the equilibrium profit function  $\pi_i(\cdot)$  can be treated as if it is exogenous (see p. 1892 of Doraszelski and Pakes 2007).

**State-to-state transitions** Here I present an abridged description of the state-to-state transitions; see the Appendix for details. In each period, the firms' respective updating and R&D investment decisions determine the industry state that arises in the next period. As explained above, in subperiod 1, a firm is able to enhance the quality of its product by releasing a new version, which incorporates each unit of the firm's R&D stock with some probability. (The uncertain nature of version releases is discussed further in Section 2.2.) Specifically, if firm  $i$  possesses  $\omega_i^R$  units of R&D stock and it releases a new version, its product quality improves by  $\bar{\omega}_i^R \in \{0, 1, \dots, \omega_i^R\}$  units with probability  $s(\bar{\omega}_i^R | \omega_i^R)$ . This resets firm  $i$ 's R&D stock,  $\omega_i^R$ , to zero.<sup>11</sup>

Consider, for example, the industry state (0,4,4). If only firm 1 releases a new version and it successfully incorporates three units of its R&D stock, then the industry transitions to state (3,0,4). If only firm 2 releases a new version and it successfully

<sup>10</sup>The assumption that there is no outside option has two noteworthy implications. First, the model admits an equilibrium in which each firm sets an infinite price and sells to half the market. I rule this out by assuming that prices are finite. Second, the Caplin and Nalebuff (1991) proof of existence and uniqueness does not apply. However, I have succeeded in computing a Nash equilibrium for every value of  $\omega^m$  at every parameterization explored in this paper. Moreover, I have not encountered multiple equilibria.

<sup>11</sup>I have assumed that if a firm fails to incorporate a unit of R&D stock into a new version, then that unit of R&D stock is lost. The reasoning and the examples that I provide in Section 2.2 suggest that this assumption is reasonable for CPG and consumer electronics categories, among others. Alternatively, one could assume that if a firm fails to incorporate a unit of R&D stock, it retains it and can try again in the future. This entails assuming that when a firm releases a new version, its R&D stock transitions from  $\omega_i^R$  to  $\omega_i^R - \bar{\omega}_i^R$  instead of zero; the rest of the model is unchanged. I thank Referee 1 for this suggestion.



incorporates two units of its R&D stock, then the industry transitions to state  $(-2, 4, 0)$ . Finally, if both firms 1 and 2 release new versions and they successfully incorporate three and two units of R&D stock respectively, then the industry transitions to state  $(1, 0, 0)$ .

In subperiod 2, the industry is initially in state  $\omega'$ . Firm  $i$ 's R&D stock for the subsequent period,  $\omega_i^{R''}$ , is determined by the stochastic outcome of its investment decision:

$$\omega_i^{R''} = \omega_i^{R'} + v_i,$$

where  $v_i \in \{0, 1\}$  is a random variable governed by firm  $i$ 's investment  $x_i \geq 0$ . If  $v_i = 1$ , the investment is successful and firm  $i$ 's R&D stock increases by one. The probability of success is  $\frac{\alpha x_i}{1 + \alpha x_i}$ , where  $\alpha > 0$  is a measure of the effectiveness of investment.

**Equilibrium** I restrict attention to symmetric Markov perfect equilibria in pure strategies. Theorem 1 in Doraszelski and Satterthwaite (2010) establishes that such an equilibrium exists.

## 2.2 Model discussion

In this section, I discuss the assumption that there is no outside option, and the assumption that R&D stocks are mutually observable. I then motivate the uncertainty that characterizes the version release process.

**No outside option** As explained above, by assuming that there is no outside option, I restrict attention to the business stealing effect and abstract from the market expansion effect. This same assumption is implicitly made in the related analytic theory literature that uses dynamic quality ladder models to study R&D competition.<sup>12</sup> While the papers in that literature make this assumption for the purpose of analytic tractability, I make it for the purpose of computational tractability.<sup>13</sup> Moreover, like the aforementioned papers, by restricting attention to the business stealing effect, I focus on how firms strategically use R&D to improve competitive positioning. Finally, assuming that there is no outside option may be a reasonable abstraction when considering *essential* CPG products such as diapers and toothpaste, which are discussed in Section 1, among others. It would be interesting to explore the effects of both business stealing and market expansion effects. However, because doing so would dramatically increase computational burden, I leave this for future work.

**Mutually observable R&D stocks** I assume that each firm can observe its rival's R&D stock. In some industries—e.g., software and biotechnology—it is common practice for firms to publicly announce intermediate R&D successes (Jansen 2010).

<sup>12</sup>See the papers cited in footnotes 5 and 6.

<sup>13</sup>In Section 3, I set  $L_m = 20$  and  $L_R = 10$ . It follows that the number of industry states is  $41 \times 11^2 = 4961$ . However, had I not reduced the dimensionality of the state space (from four to three), then there would have been  $21^2 \times 11^2 = 53,361$  industry states.

Furthermore, a firm might be able to gather information on a rival's undisclosed R&D activities through competitive intelligence (West 2001). That being said, in some industries, a firm might learn how successful a rival's R&D has been only once the rival releases a new version. It would certainly be interesting to explore a model in which firms *cannot* observe each other's R&D stocks, and therefore possess beliefs about them that evolve over time and take into consideration the information divulged by version releases. However, because this would significantly complicate the model, I leave this for future work.

**Uncertainty in the version release process** Before a new version is released, there are several reasons why a firm faces uncertainty as to how successful it will be. First, a firm is often uncertain about how much consumers will value product improvements.<sup>14</sup> Second, there are various reasons why a new feature might fail from a technical standpoint—e.g., incompatibility with the base product or with other new features—even after rigorous testing. Mennen's 1972 launch of a new version of its deodorant, which incorporated vitamin E, was plagued by both of these problems. First, despite an extensive advertising campaign, customers were not able to understand how vitamin E improved deodorant and therefore did not value the addition (Rivkin and Sutherland 2004). Second, many customers suffered allergic reactions and, as a result, the product had to be pulled from the market (Rietschel and Fowler 2008). (On a separate note, in Section 3, I explain how incorporating uncertainty into the version release process mitigates the effect of the edge of the state space.)

### 2.3 Model derivations

In this section, I derive the Bellman equation, discuss the optimality conditions (see the Appendix for a full derivation), and describe the system of equations that characterizes an equilibrium.

**Bellman equation** To derive the Bellman equation, I first consider the investment decisions that firms make in subperiod 2 and then the updating decisions that they make in subperiod 1. I let  $V_i(\omega, \phi_i)$  denote the expected net present value of all future cash flows to firm  $i$  in industry state  $\omega$  in subperiod 1, immediately after it has drawn release cost  $\phi_i$ . Firm  $i$ 's value function is  $V_i : \Omega \times \Phi \rightarrow \mathbb{R}$  and its policy functions  $x_i : \Omega \rightarrow \mathbb{R}$  and  $r_i : \Omega \rightarrow \mathbb{R}$  specify its R&D investment and its probability of

<sup>14</sup>First, some uncertainty exists irrespective of how much test marketing a firm does. Furthermore, CPG firms often elect to do relatively little test marketing because it is "slow, expensive, and open to spying and sabotage" (Baker et al. 2000). Second, the presence of such uncertainty relates to the idea that firms sometimes inadvertently focus on improving products from a technical standpoint—instead of focusing on satisfying customers' needs and wants—and therefore may make improvements that customers do not value (Levitt 1960). Finally, although Coke is not characterized by frequent version releases, the infamous 1985 release of a new version of Coke—which replaced the previous version—provides an excellent example of this phenomenon. The new version of Coke failed because of tremendous public backlash despite much market research suggesting that it would be a success (Prendergrast 2000).

updating in industry state  $\omega$ . (I derive the latter by integrating out the release cost  $\phi_i$ .) As I explain further below, because I solve for a *symmetric* equilibrium, it suffices to solve for the optimality conditions for one firm; hence I hereafter restrict attention to firm 1.

*Investment decision* At the beginning of subperiod 2, the industry is in state  $\omega'$ . The expected net present value of cash flows to firm 1 is

$$U_1(\omega') \equiv \max_{x_1 \geq 0} \left\{ -x_1 + \beta E \left[ V_1(\omega'', \phi'_1) | \omega', x_1 \right] \right\},$$

where  $\beta \in (0, 1)$  is the discount factor. Firm 1 chooses investment  $x_1 \geq 0$  that maximizes the expected net present value of its future cash flows. In the [Appendix](#), I derive the optimality condition for firm 1's R&D investment,  $x_1(\omega')$ ; see [Appendix Eq. 3](#).

*Version release decision* At the beginning of subperiod 1, the industry is in state  $\omega$ . Let firm 1's perceived probability that firm 2 releases a new version of its product be  $r_2(\omega)$ . (In equilibrium,  $r_2(\omega)$  will be determined by firm 2's equilibrium updating strategy.) Firm 1's value function  $V_1 : \Omega \times \Phi \rightarrow \mathbb{R}$  is implicitly defined by the Bellman equation

$$V_1(\omega, \phi_1) = \pi_1(\omega^m) + \max_{\chi \in \{0,1\}} (1 - \chi) \left\{ (1 - r_2(\omega))U_1(\omega) + r_2(\omega)Y_1^2(\omega) \right\} + \chi \left\{ -\phi_1 + (1 - r_2(\omega))Y_1^1(\omega) + r_2(\omega)Y_1^{12}(\omega) \right\}, \tag{1}$$

where  $Y_1^1(\omega)$  is the expected net present value of firm 1's future cash flows in industry state  $\omega$  if it decides to update and firm 2 does not;  $Y_1^2(\omega)$  is defined analogously for the case in which firm 1 does not update and firm 2 does; and  $Y_1^{12}(\omega)$  is defined analogously for the case in which both firms update.

In the optimization problem on the right-hand side of Bellman Eq. 1, firm 1 releases a new version ( $\chi = 1$ ) only if the new version yields greater value than its existing product, net of the release cost  $\phi_1$ . It follows that firm 1 plays a threshold updating strategy; i.e., if it draws a sufficiently low release cost  $\phi_1$ , it releases a new version and otherwise it does not. I derive this updating threshold and then, by integrating over the release cost  $\phi_1$ , I derive a closed-form expression for firm 1's optimal probability of updating,  $r_1(\omega)$ ; see [Appendix Eq. 8](#).

**Solving for an equilibrium** Because I solve for a *symmetric* Markov perfect equilibrium, the investment decision taken by firm 2 in state  $\omega$  is identical to the investment decision taken by firm 1 in state  $\omega^{[2]} \equiv (-\omega^m, \omega_2^R, \omega_1^R)$ , i.e.,  $x_2(\omega) = x_1(\omega^{[2]})$ . A similar relationship holds for the probability of releasing a new version and the value function:  $r_2(\omega) = r_1(\omega^{[2]})$  and  $V_2(\omega, \phi_2) = V_1(\omega^{[2]}, \phi_2)$ . It therefore suffices

to determine the value and policy functions of firm 1. Solving for an equilibrium for a particular parameterization of the model amounts to finding a value function  $V_1(\cdot)$  and policy functions  $x_1(\cdot)$  and  $r_1(\cdot)$  that satisfy the Bellman equation and the R&D investment and updating optimality conditions in [Appendix Eqs. 3, 8 and 9](#) respectively for all industry states  $\omega \in \Omega$ .

### 3 Computation

In this section, I explain how incorporating uncertainty into the version release process mitigates the effect of the edge of the state space. I also present the baseline parameterization and discuss the algorithm that I use to compute equilibria.

**Mitigating the effect of the edge of the state space** The incentives that a firm faces when its R&D stock reaches the maximal level ( $\omega_i^R = L^R$ ) are different from those it faces elsewhere because it cannot further increase its R&D stock. This general issue—that the edge of the state space distorts incentives—arises commonly in models in the Ericson and Pakes (1995) framework. It is typically addressed by assuming that diminishing returns set in sufficiently quickly as a firm approaches the edge of the state space.<sup>15</sup> It follows that from a firm’s perspective, the states on the edge are not too dissimilar from the states near the edge, which mitigates the effect of the edge of the state space.

I address this issue by introducing diminishing returns to R&D stock accumulation. Recall that when a firm releases a new version, it successfully incorporates each unit of its R&D stock with some probability; I assume that each additional unit of R&D stock is less likely to be successfully incorporated than the previous unit. Specifically, I assume that the  $n^{\text{th}}$  unit of a firm’s R&D stock is successfully incorporated with probability  $\rho^n$ , where  $\rho \in (0, 1)$ . It follows that when a firm attempts to incorporate a large R&D stock (say  $\omega_i^R = 10$ ) into a new product, it is very unlikely to successfully incorporate all of its R&D stock units. See the [Appendix](#) for further detail.

Recall that a firm is uncertain as to how successful a new version of its product will be because (i) consumers may not value the improvements that are incorporated into the product, and (ii) the “improvements” may fail to be successfully incorporated for technical reasons. Therefore, the decreasing returns described above can be interpreted in two ways. First, the greater the number of new features that a firm incorporates into a new version, the less likely a consumer is to value each additional feature. This reflects the notion that consumers have limited attention for product attributes (Dahremoller and Fels 2015). Second, the greater the number of features that a firm *attempts* to incorporate, the less likely it is to succeed in incorporating

<sup>15</sup>For example, Pakes and McGuire (1994), Gowrisankaran and Town (1997), and Borkovsky et al. (2012) assume that a consumer’s utility is characterized by diminishing returns to quality that set in very quickly beyond some quality threshold. I am unable to take this particular approach because in order to reduce the dimensionality of the state space, I have assumed that a consumer’s utility is linear in quality.

**Table 1** Baseline parameterization

Parameter	$L^m$	$L^R$	$M$	$c$	$b$	$\sigma$	$\alpha$	$\rho$	$G_l$	$G_u$	$\beta$
Value	20	10	7	5	0.3	0.1	1	0.95	0	80	0.951


each additional feature. This reflects the notion that it is more difficult to successfully incorporate many new features into a new version than it is to successfully incorporate only a few.

**Baseline parameterization** I provide a few comments on the baseline parameterization, which is presented in Table 1. First, as  $\frac{L^m}{L^R} = 2$ , a firm would have to make *at least* two version releases (while its rival made no releases whatsoever) in order to move from a state at which firms are tied in the product market to a state at which it achieves the maximal product market lead of 20 units. Selecting a relatively large  $L^m$  value ensures that the edges of the state space at which  $\omega^m = 20$  and  $\omega^m = -20$  do not adversely affect equilibrium behavior. This is because the leader/follower roles that arise in equilibrium are established long before either firm begins to approach a state in which it achieves the maximal product market lead.<sup>16</sup>

Second, the release cost,  $\phi_i$ , includes all costs that a firm might incur as it strives to transform its current product and the innovations that it has achieved since launching it into a new commercially-viable product that is successfully brought to market. This will typically include the costs required to incorporate the new innovations into the core product as well as the costs of a product launch and related marketing activities. However, the release cost also accounts for many possible complications. First, incorporating the new innovations might require unanticipated changes to the material requirements and/or the production process. Second, the new product might necessitate modifications to the existing packaging. Third, the existing distribution channel might not be ideally suited to the new product. Fourth, an important retailer (e.g., Walmart) might not immediately agree to stock the new product, or it might request product modifications. Finally, it might be impossible to simultaneously satisfy the different product specifications of two important retailers (e.g., Walmart and Target). Because the release cost incorporates a wide variety of different costs that *might* be incurred in launching a new version, I assume that the release cost is highly uncertain—specifically, that it is drawn from a uniform distribution over the interval  $[G_l, G_u] = [0, 80]$ .<sup>17</sup>

<sup>16</sup>In the [Appendix](#), I show that a version of the model with a very small state space is not rich enough to admit the results presented in Section 6.

<sup>17</sup>Additionally, allowing for a wide range of possible release costs yields equilibria with a wide range of release probabilities (across industry states), which makes it easier to discern the different equilibrium behaviors that arise. I have verified that qualitatively similar behaviors arise for narrower ranges of release costs. However, narrower ranges of release costs tend to give rise to greater convergence problems for the algorithm described below. Finally, Section 6.2 explores the effects of changes in the range of release costs by computing equilibria for higher values of  $G_l$  and  $G_u$ .



**Fig. 1** Equilibrium price  $p_1(\omega^m)$  (left panel), profit  $\pi_1(\omega^m)$  (middle panel), and the (discrete) rate of increase of profit  $\pi_1(\omega^m) - \pi_1(\omega^m - 1)$  (right panel) for  $\sigma = 0.1, 2,$  and  $4$

Third, the discount rate of  $\beta = 0.951$  corresponds to a period length of four months, an interest rate of 5% on an annualized basis, and an expected industry lifespan of 10 years—that is, in each period, the industry dies with probability  $\frac{1}{10 \times \frac{12}{4}} = \frac{1}{30}$ .

**Algorithm** I first solve for the Nash equilibrium of the product market game (for each product market state  $\omega^m$ ) by numerically solving the system of optimality conditions corresponding to the firms’ respective profit-maximization problems. The equilibrium profit function  $\pi_i(\cdot)$  is then treated as an input to the (Pakes and McGuire 1994) algorithm, which is used to compute Markov perfect equilibria of the dynamic model. To explore the equilibrium correspondence, I nest the (Pakes and McGuire 1994) algorithm in a simple continuation method (Judd 1998); see the Appendix for details.

#### 4 Results: product market competition

Before exploring the impact of changes in the the degree of horizontal differentiation on firms’ (dynamic) product development strategies, I explore the implications of such changes for firms’ (static) pricing strategies. In this section, I present the equilibria of the (static) product market game for low ( $\sigma = 0.1$ ), moderate ( $\sigma = 2$ ), and high ( $\sigma = 4$ ) degrees of horizontal differentiation. Firm 1’s equilibrium price and profit functions are presented in Fig. 1, and its equilibrium market share functions are presented in the Appendix.<sup>18</sup>

<sup>18</sup>Because firms are symmetric, firm 2’s equilibrium price and profit functions are symmetric to firm 1’s, i.e.,  $p_2(\omega_m) = p_1(-\omega_m)$  and  $\pi_2(\omega_m) = \pi_1(-\omega_m)$ .

**A low degree of horizontal differentiation** A low degree of horizontal differentiation ( $\sigma = 0.1$ ) has several important implications. First, firms engage in aggressive price competition; hence, when they are tied in the product market—because neither firm has a quality advantage that it can leverage—they both earn very low profits. Second, a firm that gains even a small lead dominates the market in terms of share. Moreover, its profits increase relatively rapidly in the size of its lead because the larger its quality advantage, the higher the price it charges. Third, because a quality laggard sets very low prices and commands little market share, it earns extremely low profits that are highly unresponsive to the size of the leader’s lead.

**A moderate degree of horizontal differentiation** When the degree of horizontal differentiation is moderate ( $\sigma = 2$ ), as when it is low, a firm’s price and profits are strictly increasing in the size of its lead. However, due to the higher degree of horizontal differentiation, firms engage in less aggressive price competition, and a quality advantage does not translate into as large an increase in market share. This has several implications. First, the leader’s price and market share—and accordingly its profits—do not increase as rapidly as its lead grows. Second, a follower earns sizeable profits that increase substantially as it narrows the leader’s lead. So, unlike the  $\sigma = 0.1$  scenario, a follower benefits from narrowing the leader’s lead even if it never overtakes the leader, or from slowing the rate at which the leader expands its lead. Third, following from the first two points, the benefits of leadership in the product market, while still significant, are not as great as when the degree of horizontal differentiation is low.

The first two points are summarized in the the right panel of Fig. 1, which presents the (discrete) rate at which a firm’s profits increase as its quality advantage grows (or its quality disadvantage declines). It emphasizes that an increase in the degree of horizontal differentiation makes the profit function flatter for the quality leader, but steeper for the follower; hence, it *weakens* the leader’s short-run incentives to invest in product development, but it *strengthens* those of the follower. The asymmetric impact of an increase in the degree of horizontal differentiation on the (short-run) investment incentives of the leader and follower has important implications, which are explored in Section 6.1.

**A high degree of horizontal differentiation** A high degree of horizontal differentiation ( $\sigma = 4$ ) induces firms to engage in even less aggressive price competition, and it only further softens the impact of a quality advantage on market share. Hence, it further reduces the benefits of leadership in the product market. Moreover, as the right panel of Fig. 1 emphasizes, the profit function becomes even flatter for the leader, and even steeper for the follower. Hence, the leader’s (short-run) incentives to invest in product development become even weaker, and the follower’s become even stronger.

## 5 A benchmark model with no version releases

Before turning to the results of the dynamic model presented in Section 2 (hereafter, the “full model”), I present results for a simplified version of the model (nested in

the full model) in which firms cannot stockpile R&D and accordingly do not make version release decisions. Specifically, I assume that when a firm achieves an R&D success, it is immediately incorporated into its product with certainty and at no additional cost.<sup>19,20</sup> Benchmarking the model in Section 2 against this model will allow me to show exactly how the ability to stockpile R&D and decide on the timing of version releases impacts firms' product development strategies.

In the benchmark model, because firms cannot stockpile R&D, the industry state is simply the product market state  $\omega^m$ . Because firms do not make version release decisions, each firm has only an R&D investment policy function. The model of product market competition is unchanged. In the Appendix, I restate the dynamic model while incorporating the simplifying assumptions described above.

In Fig. 2, I present equilibria for low ( $\sigma = 0.1$ ), moderate ( $\sigma = 2$ ), and high ( $\sigma = 4$ ) degrees of horizontal differentiation. (I found only one equilibrium for each parameterization.) Firm 1's R&D investment policy function is presented in the left panels.<sup>21</sup> To explore the evolution of industry structure over time, I use the firms' equilibrium policy functions to compute the transient distribution over industry states in period  $t$  starting from state  $\omega^m = 0$  in period 0. In the right panels of Fig. 2, I present the transient distributions for  $t = 30$ , which is the industry's expected lifespan, and  $t = 100$ , which represents the long run.

The upper-left panel of Fig. 2 shows that when the degree of horizontal differentiation is low ( $\sigma = 0.1$ ), firms engage in R&D preemption races; i.e., when neither firm has a large lead, each firm invests heavily in R&D in an attempt to become the market leader. Once one firm gains a lead of three units, it induces its rival to "give up"—i.e., to cease investing altogether; thereafter, because the rival no longer poses a threat to the leader's leadership status, the leader reduces its R&D investment. The upper-right panel of Fig. 2 shows that this behavior gives rise to an asymmetric industry structure in the short and long run.

Preemption races are driven by the benefits of leadership in the product market. As explained in Section 4, because the degree of horizontal product differentiation is very low, a firm that gains only a small lead dominates the product market. It follows that each firm faces very strong incentive to become and remain the quality leader. Hence, when firms are tied in the product market, each firm invests heavily in R&D in an attempt to become the quality leader.

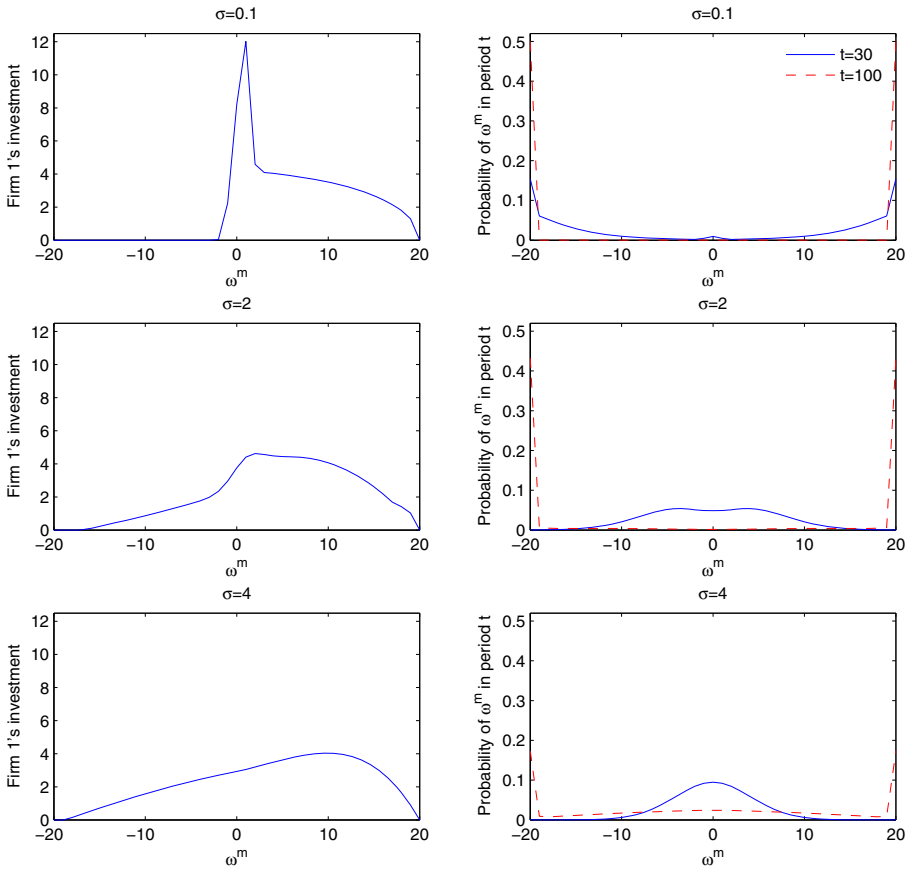
Moreover, in the middle panel of Fig. 1, one can see that the profit function is extremely flat for the follower ( $\omega_1 < 0$ ) and is quite steep for the leader ( $\omega_1 > 0$ ); hence, in the short run the follower can increase its profit negligibly by reducing the size of the leader's lead, but the leader can increase its profit significantly by

<sup>19</sup>In this respect, the simplified model is similar to earlier quality ladder models (e.g., Pakes and McGuire 1994, Borkovsky et al. 2012), the only difference being that, as explained in Section 2.1, I reduce the dimensionality of the state space.

<sup>20</sup>One derives the simplified model from the model in Section 2 by setting the release cost to zero ( $G_i = G_u = 0$ ) and the probability that a firm successfully incorporates its entire R&D stock (when releasing a new version) to one—i.e.,  $s(\omega_i^R|\omega_i^R) = 1$ .

<sup>21</sup>Because firms are symmetric, firm 2's R&D investment policy function is symmetric to firm 1's, i.e.,  $x_2(\omega^m) = x_1(-\omega^m)$ .





**Fig. 2** Model with no version releases. Equilibria for  $\sigma = 0.1$  (top row),  $\sigma = 2$  (middle row), and  $\sigma = 4$  (bottom row). R&D investment policy function  $x_1(\omega^m)$  (left panels) and transient distributions  $\mu_t^m(\omega^m)$  for  $t = 30$  and  $t = 100$  (right panels)

expanding its lead. It follows that a leader faces much stronger (short-run) incentive to invest in R&D than a follower. In a preemption race, each firm tries to exploit these incentives. That is, a firm invests heavily in R&D in hopes of gaining a lead over its rival because that lead would strengthen its own incentive to innovate and weaken its rival's incentive to innovate. Accordingly, when one firm does gain a lead, it invests even more heavily and its rival invests less heavily, making it likely that the lead will only grow. This continues until the lead is large enough to induce the follower to give up, explaining why an asymmetric industry structure arises.

The middle-left panel of Fig. 2 shows that at a moderate level of horizontal differentiation ( $\sigma = 2$ ), the preemption race is much milder, for several reasons. First, because the benefits of leadership are lower (as explained in Section 4), firms do not compete as aggressively for the lead. Second, because the higher degree of horizontal differentiation strengthens the follower's (short-run) investment incentives and weakens those of the leader (as explained in Section 4), a lead is not as beneficial in

terms of its differential impact on the investment incentives of leader and follower. Finally, for this same reason, the leader needs to achieve a *much* larger lead to induce the follower to give up, and this too weakens firms' incentives to engage in preemption. As shown in the middle-right panel, it follows that the industry structure is more symmetric in both the short and long run.

The bottom panels of Fig. 2 shows that at a high level of horizontal differentiation ( $\sigma = 4$ ), due to an amplification of the effects described above, firms engage in no preemption whatsoever, and this leads to a more symmetric industry structure in both the short and long run.

## 6 Results: the dynamic model of product development

In this section, I present results for the dynamic model of product development described in Section 2. Specifically, I explore the effects of two key industry fundamentals—the degree of horizontal product differentiation and the cost of releasing a new version—on product development strategies and, accordingly, the evolution of industry structure.

### 6.1 Horizontal differentiation

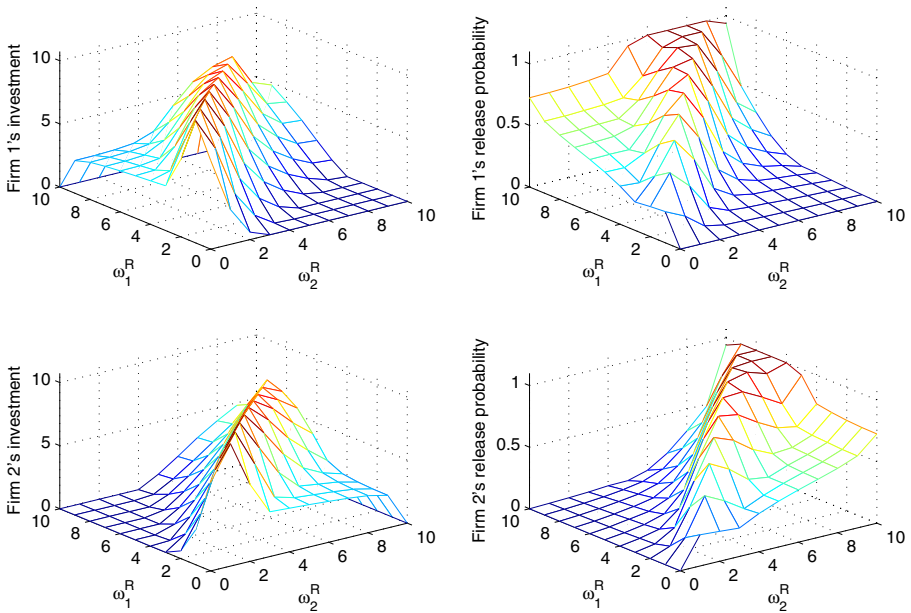
In this section, I show that low, moderate, and high degrees of horizontal differentiation give rise to three distinctly different approaches to product development.<sup>22</sup> Furthermore, by comparing the results to those of the benchmark model, I show that two of these approaches—for moderate and high degrees of horizontal differentiation—arise only if firms can stockpile R&D and decide exactly when to release new versions.

#### 6.1.1 A low degree of horizontal differentiation

I first present the equilibrium for the baseline parameterization, at which the degree of horizontal differentiation is low ( $\sigma = 0.1$ ). The equilibrium investment and probability of updating functions are mappings from the three-dimensional state space  $\Omega$  to  $\mathbb{R}$ ; therefore, these functions are four-dimensional. I present equilibria by graphing three-dimensional cross-sections of these functions. Specifically, holding firm 1's product market lead  $\omega^m$  fixed, I graph  $x_1(\omega^m, \cdot, \cdot)$ ,  $r_1(\omega^m, \cdot, \cdot)$ ,  $x_2(\omega^m, \cdot, \cdot)$  and  $r_2(\omega^m, \cdot, \cdot)$ . Throughout the paper, I present cross-sections that best illustrate the equilibrium behaviors that arise. (All cross-sections of all equilibria presented and summarized below are available upon request.)

The  $\omega^m = 0$  cross-sections of the equilibrium investment and probability of updating functions for both firms are presented in Fig. 3. From these figures, one can see how firms behave when they are tied in the product market. In particular, these figures

<sup>22</sup>I also compute equilibria for fine discretization of a wide range of  $\sigma$  values and present the equilibrium correspondence in the [Appendix](#).



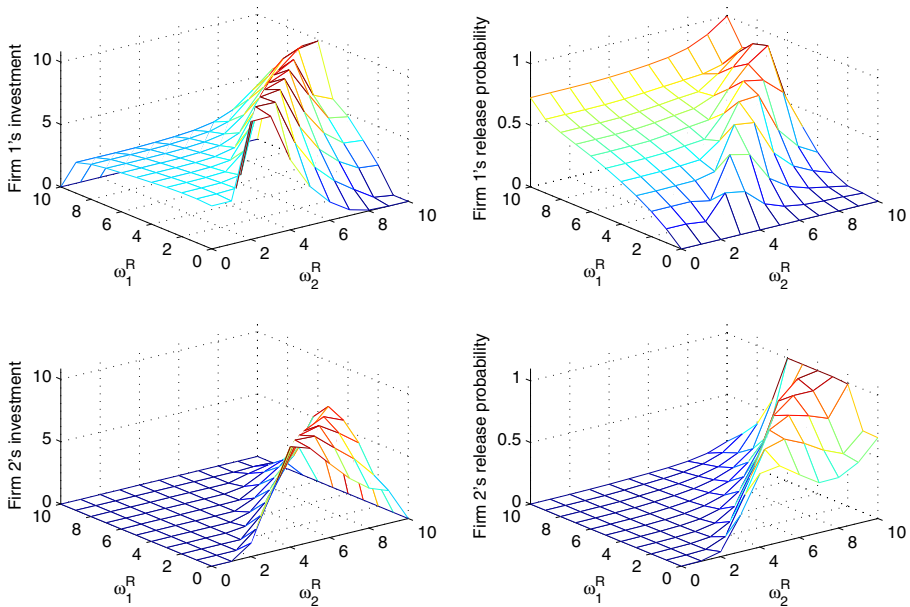
**Fig. 3** Preemption ( $\sigma = 0.1$ ).  $\omega^m = 0$  cross-sections of policy functions for R&D investment  $x_i(\omega)$  and release probability  $r_i(\omega)$  for firm 1 (*top panels*) and firm 2 (*bottom panels*)

present each firm’s investment choice and probability of updating for any possible pair of R&D stock levels  $(\omega_1^R, \omega_2^R)$ . Firm 1’s equilibrium investment and probability of updating functions are in the top row of the figure; those of firm 2 are in the bottom row.<sup>23</sup>

**Preemption races** Figure 3 shows that when firms are tied in the product market and neither firm has a large R&D stock lead, firms engage in a preemption race in which they invest heavily in R&D and release new versions with high probability. (Hence, I refer to this equilibrium as the *preemptive equilibrium*.) This behavior is reflected in the pronounced ridges—which I call *preemptive ridges*—that lie just off the diagonals of the cross-sections of all four policy functions. Figure 3 also shows that a firm can induce its rival to “give up”—i.e., to cease investing altogether and to release new versions with very low probability—by gaining a sufficiently large R&D stock advantage.<sup>24</sup> Hence, the preemption race can end even if neither firm has

<sup>23</sup>In this case, the matrices graphed in the bottom row are transposes of the corresponding matrices in the top row; this is because I restrict attention to symmetric equilibria and because firms are tied in the product market. However, this will not be the case for cross-sections for which  $\omega^m \neq 0$ .

<sup>24</sup>The bottom panels in Fig. 3 show that there is a roughly triangular region in the  $(\omega_1^R, \omega_2^R)$  grid—where  $\omega_1^R$  is sufficiently high and  $\omega_2^R$  is sufficiently low—in which firm 2 gives up. The upper panels show that there is an analogous region in which firm 1 gives up.



**Fig. 4** Preemption ( $\sigma = 0.1$ ).  $\omega^m = 3$  cross-sections of policy functions for R&D investment  $x_i(\omega)$  and release probability  $r_i(\omega)$  for firm 1 (top panels) and firm 2 (bottom panels)

released a new version and, accordingly, firms are still tied in the product market. Once a rival has given up, it ceases to pose a significant threat to the leader, and the leader responds by significantly reducing its R&D investment and probability of updating. Before explaining why this behavior arises—and, specifically, how it relates to the preemption race of the benchmark model presented in Section 5—I present another cross-section of this equilibrium and then discuss the evolution of industry structure over time.

Figure 4 presents the  $\omega^m = 3$  cross-sections of the equilibrium policy functions. From these figures, one can see how firms behave when firm 1 has a lead of three in the product market. These cross-sections are qualitatively similar to the  $\omega^m = 0$  cross-sections. However, the preemption race does *not* occur when firms have similar R&D stocks—as in Fig. 3—but rather when firm 2's R&D stock lead is large enough to threaten firm 1's lead in the product market. In the heat of a preemption race, if both firms simultaneously release new versions, firm 2 is likely to catch up to firm 1 in the product market.

Cross-sections of the equilibrium for other  $\omega^m \in \{-10, \dots, 10\}$  are qualitatively similar to the cross-sections presented above; a preemption race occurs when the product market follower has a large enough R&D stock lead such that the *expected* product market lead would be close to zero if both firms updated simultaneously. Once the product market leader achieves a lead of eleven or more units, the follower ceases to invest; hence, firms no longer engage in preemption races. The qualitative differences between different cross-sections of the equilibria presented later in the

paper are similar in nature to those discussed above. Therefore, I hereafter present only the  $\omega^m = 0$  cross-sections of equilibrium policy functions.

To explore the evolution of industry structure over time, I use the equilibrium policy functions to compute  $\mu_t(\cdot)$ , the transient distribution over states in period  $t$  starting from state  $(0, 0, 0)$  in period 0. I then compute the transient distribution over product market states

$$\mu_t^m(\omega^m) = \sum_{\omega_1^R=0}^{10} \sum_{\omega_2^R=0}^{10} \mu_t(\omega^m, \omega_1^R, \omega_2^R)$$

for  $\omega^m \in \Omega_d^m$ , and I use this to compute the size of the leader’s expected lead in period  $t$

$$L^t = \sum_{\omega^m=-20}^{20} \mu_t^m(\omega^m) \times |\omega^m|.$$

In the left panel of Fig. 6, I present the transient distributions over product market states for  $t = 30$  and  $t = 100$ . The industry is quite asymmetric after 30 periods, and extremely asymmetric in the long run; the size of the leader’s expected lead after 30 (100) periods is 15.94 (20.00).<sup>25</sup>

**Why do firms engage in preemption races?** The incentives that give rise to the preemption races observed in Figs. 3 and 4 bear some similarity to those that underpin the preemption race of the benchmark model. Both the full model and the benchmark model include the same static model of product market competition; therefore in both, when the degree of horizontal differentiation is low ( $\sigma = 0.1$ ), firms compete aggressively for the lead because the benefits of leadership in the product market are high and a firm can dominate it by gaining only a small lead.

That being said, there are some features of the preemption races in Figs. 3 and 4 that are not explained by the benchmark model. First, the left panels of Fig. 3 show that even if the firms remain tied in the product market, a firm can influence its rival’s incentive to innovate by building an R&D stock advantage—inducing its rival to invest less in R&D and, if it falls sufficiently far behind, even to cease investing altogether. This means that a firm can achieve the leadership position in the indus-

<sup>25</sup>When one firm gains a lead of 20 in the product market, both firms neither invest nor release new versions, i.e., the industry reaches an absorbing state. It does however take a very long time until this occurs—approximately 60 periods, which is twice the industry’s expected lifespan. The standard approach to addressing this issue in Ericson and Pakes (1995) models is to assume that a firm experiences decreasing returns as it approaches the upper edge of the state space. I have verified that a version of the model in which the effectiveness of investment,  $\alpha$ , decreases in the size of a firm’s lead yields equilibria that are qualitatively similar to those presented in Section 6, the only difference being that firms never stop investing or releasing new versions and accordingly neither firm ever achieves a maximal product market lead of 20. These equilibria are available upon request.

try without releasing even a single new version. This raises the question of why the rival would reduce its R&D investment and even give up altogether if it has not fallen behind in the product market. The reason is that firms' R&D investment incentives are shaped not only by the current product market state, but also by their expectations over future product market states. Firms understand that—given the equilibrium updating probabilities—a firm with an R&D stock advantage is likely to soon become the product market leader.<sup>26</sup> Accordingly, a firm that finds itself at an R&D stock disadvantage anticipates that it will soon become the product market laggard and—because the profit function for the product market laggard is so flat—it faces extremely weak R&D investment incentives.

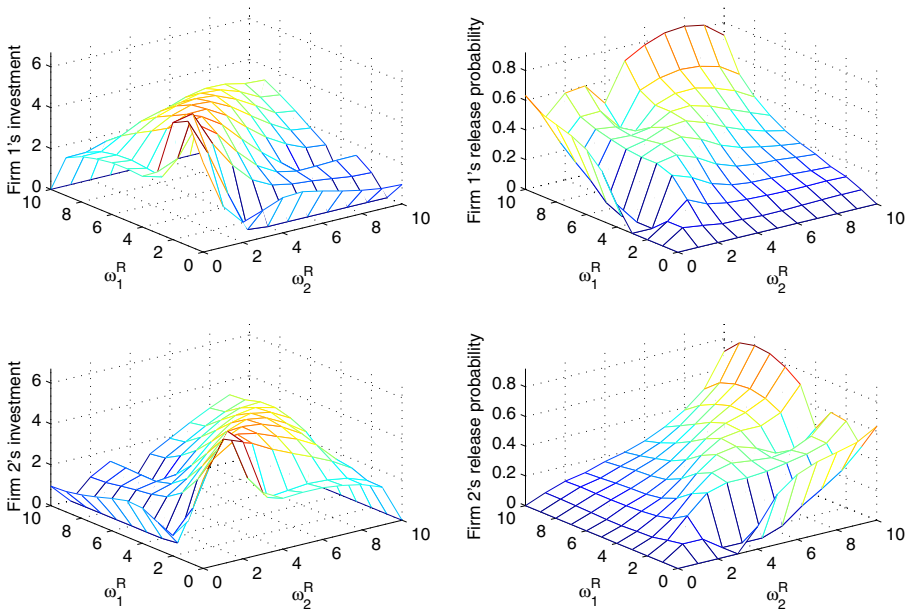
The above discussion emphasizes the strategic nature of preemption and, specifically, the strategic advantage that stems from having an R&D stock advantage. However, this raises an important question: In a preemption race, firms not only invest heavily in R&D, but also release new versions with high probability. It follows that a version release must *also* confer a strategic advantage.<sup>27</sup> However, while a version release can certainly enhance (static) profits, if it simply entails transferring R&D successes from a firm's R&D stock into its product, why does it also generate a strategic advantage? In other words, why would it impact firms' incentives to innovate?

Consider the industry state  $(0, 1, 1)$ —in which firms are tied in the product market and in terms of R&D stock. Suppose that firm 1 releases a new version and successfully incorporates its unit of R&D stock, bringing about a transition to industry state  $(1, 0, 1)$ . While firm 1 now has a one unit lead in the product market, it is one unit behind in the R&D stock race. So, it might seem as if the firms are still effectively tied, which suggests that the version release should have no *strategic* impact whatsoever. However, the version release does indeed have a strategic impact, and it stems directly from the fact that the version release process is uncertain. In industry state  $(1, 0, 1)$ , firm 2 finds itself at a disadvantage because to catch up to firm 1 in the product market it would have to release a new version *and* successfully incorporate its unit of R&D stock; however, there is no guarantee that the latter will occur. This disadvantage weakens firm 2's incentive to innovate. In the [Appendix](#), I show that if the uncertainty in the version release process is “turned off”, then firms engage in much less aggressive preemption in terms of version releases.<sup>28</sup>

<sup>26</sup>For example, in industry state  $(0, 6, 2)$ —in which firm 2 gives up—firm 1 is the likely product market leader not only because of its R&D stock advantage, but also because it updates with much higher probability (0.4272) than firm 2 (0.0395).

<sup>27</sup>A version release has a direct effect and a strategic effect. The direct effect is that a version release enhances a firm's product quality and accordingly its profits. The larger a firm's R&D stock, the more it enhances the firm's (expected) profits when the firm releases a new version, and accordingly the higher the firm's release probability should be. Hence, if one considers only the direct effect, one would expect a firm's release probability to be strictly increasing in its R&D stock, as is the case in a monopolistic version of the model. The fact that a firm's release probability is *not* strictly increasing in its R&D stock can be attributed to the strategic effect, i.e., the impact of a version release on a rival firm's behavior.

<sup>28</sup>Some preemption in terms of version releases still remains because of the uncertain nature of the release cost; see the [Appendix](#) for details.



**Fig. 5** Accommodation ( $\sigma = 2$ ).  $\omega^m = 0$  cross-sections of policy functions for R&D investment  $x_i(\omega)$  and release probability  $r_i(\omega)$  for firm 1 (top panels) and firm 2 (bottom panels)

The discussion above emphasizes that firms engage in preemption in terms of both R&D investment and version releases because each can confer a strategic advantage. While R&D preemption arises in multistage patent races<sup>29</sup> and in quality ladder models (Borkovsky et al. 2012), this paper shows that firms *also* preempt in terms of version releases and helps us understand why. (More importantly, none of the earlier models admit the behaviors that are described in Sections 6.1.2 and 6.1.3.)

### 6.1.2 A moderate degree of horizontal differentiation

I now explore the effect of a moderate degree of horizontal product differentiation ( $\sigma = 2$ ). Recall that in the benchmark model, firms' product development strategies when the degree of horizontal differentiation is moderate are qualitatively similar to those that arise when it is low; the only difference is that when it is moderate, firms engage in a milder preemption race. Here I show that the full model gives rise to a qualitatively different result, which arises because of the interaction between firms' R&D investment strategies and their version release strategies.

<sup>29</sup>In multistage patent races, if R&D investment has a deterministic impact on innovation (Fudenberg et al. 1983; Harris and Vickers 1985; Lippman and McCardle 1988), then  $\varepsilon$ -preemption arises, i.e., once a firm gains an arbitrarily small lead, it induces its rival to immediately drop out of the race. If R&D investment has a stochastic impact on innovation (Grossman and Shapiro 1987; Harris and Vickers 1987; Lippman and McCardle 1987), then the leader invests more than the follower and is therefore likely to expand its lead until it ultimately induces its rival to drop out.

Figure 5 presents the  $\omega^m = 0$  cross-sections of the equilibrium policy functions. Before turning to the trenches that appear prominently in all four panels, I will describe other aspects of the policy functions. First, because an increase in the degree of horizontal differentiation reduces the benefits of leadership in the product market, it follows that—as in the benchmark model—firms do not compete as aggressively for the lead. Hence, while firms still engage in preemption races, they are characterized by lower R&D investments and updating probabilities than in the preemptive equilibrium. Second, as in the preemptive equilibrium, a leader tends to invest more in R&D and update with higher probability than a follower. This is because a leader's short-run investment incentives are stronger than those of the follower, and because, as before, the follower determines that it would be too costly to invest in catching up. However, even though the follower accepts its role as the laggard, it does not give up altogether because, as explained in Section 4, it benefits from slowing the rate at which the leader increases its lead.

**Phases of accommodation** In Fig. 5, in the cross-sections of firm 1's policy functions in the upper panels and firm 2's investment policy function in the lower-left panel, there are two trenches: one is at  $\omega_2^R = 1$  for  $\omega_1^R \geq 3$ ; the other is at  $\omega_2^R = 4$  for  $\omega_1^R \geq 8$ .<sup>30</sup> I refer to these as *accommodative trenches*; when the industry enters such a trench, it begins a *phase of accommodation*. (I hereafter refer to this equilibrium as the *accommodative equilibrium*.<sup>31</sup>)

Consider the accommodative trench along  $\omega_2^R = 1$  for  $\omega_1^R \geq 3$ . This trench arises only once firm 1 has achieved a lead large enough to establish itself as the industry leader; once the industry reaches state  $\omega = (0, 3, 1)$ , it is extremely unlikely that firm 2 will succeed in overtaking firm 1 in the product market.<sup>32</sup> In this trench, firms accommodate each other as follows. Firm 1 invests less in R&D and releases a new version with lower probability. Hence, it slows its march toward industry dominance by both slowing the rate at which it stockpiles R&D and exercising greater patience in making its updating decisions. Firm 2 invests very little in R&D; therefore, it is unlikely to achieve an R&D success and accordingly its R&D stock is likely to remain  $\omega_2^R = 1$ .<sup>33</sup> As a result, firm 2 ceases to threaten firm 1's dominance.

Both firms benefit from this accommodation. Because firm 2 ceases to pose a threat to firm 1's leadership, firm 1 can spend less on R&D and save on the cost of

<sup>30</sup>Similarly, because this is a symmetric equilibrium and in the  $\omega^m = 0$  cross-section, firms are tied in the product market, there are identical trenches along  $\omega_1^R = 1$  for  $\omega_2^R \geq 3$  and along  $\omega_1^R = 4$  for  $\omega_2^R \geq 8$ , in which the roles of the firms are reversed.

<sup>31</sup>I have found three additional equilibria for the  $\sigma = 2$  parameterization that are qualitatively similar to the one presented in Fig. 5, the only difference being that the accommodative trenches arise at different R&D stock levels; see the Appendix for details.

<sup>32</sup>The probability that firm 2 is the product market leader ( $\omega_m < 0$ ) 100 periods after the industry enters state  $\omega = (0, 3, 1)$  is 1.88%.

<sup>33</sup>In fact, for  $\omega_1^R \geq 5$ , firm 2 does not invest at all while in this trench. Therefore, its R&D stock remains fixed at  $\omega_2^R = 1$  as long as long as neither firm releases a new version.




releasing a new version by patiently waiting for a low release cost. Firm 2 benefits because, having accepted its position as the eventual product market laggard, it is able to slow its rival's march toward industry dominance; thus it earns higher transient profits as it recedes in the product market more slowly.

I next explain how accommodation is enforced. In particular, I explain how a leader induces a follower to sharply reduce its investment. In doing so, I highlight the strategic role of R&D stockpiling—specifically, the leader's ability to influence the follower's incentive to innovate by threatening to release a new version, which would incorporate its R&D stock. Consider firm 1's policy functions in the upper panels of Fig. 5 and in particular the accommodative trench at  $\omega_2^R = 1$ . To induce firm 2 to depress its investment in this trench, firm 1 sets sufficiently high R&D investment levels and updating probabilities for the neighboring states along  $\omega_2^R = 2$ . It follows that were firm 2 to achieve an investment success, it would be punished by firm 1, which would dramatically increase its R&D investment and updating probability. Firm 2 would both become a product market laggard sooner rather than later and recede in the product market more quickly, and its profits would decrease accordingly. The results of the benchmark model in Section 5 show that accommodation does not arise in an otherwise identical model *without* R&D stockpiling and endogenous version releases; this demonstrates that the ability to use an R&D stockpile as a threat is essential for sustaining accommodation.

A phase of accommodation can come to an end in two ways. First, if the follower achieves an investment success, the industry transitions from the accommodative trench to an adjacent state that is outside of the trench. Second, if either firm elects to release a new version, the industry may transition to a different cross-section of the state space characterized by a different  $\omega^m$  value. In this other cross-section, the industry might immediately find itself in a different accommodative trench, or in a preemption race, or in neither; this all depends on the success of the version release and, accordingly, the magnitude of the leader's product market lead after the release occurs.

The right panel of Fig. 6 presents the transient distributions over product market states. It shows that after 30 (100) periods, the industry structure is asymmetric; the size of the leader's expected lead is 9.84 (19.98). Due to the albeit mild preemption races, and because the leader invests more in product development than the follower, one firm eventually gains as large a lead as is possible; however, it takes much longer under the accommodative equilibrium (90 periods) than under the preemptive equilibrium (60 periods), for several reasons. First, as explained above, the preemption races are milder. Second, the disparity between the leader's and follower's investments in product development is smaller. While these differences alone would slow the rate at which the firms diverge, the firms slow it even further by entering mutually beneficial phases of accommodation, allowing the leader to save on investment and upgrade costs and the follower to earn higher profits as it falls back more slowly than it otherwise would.

I conclude this section by explaining that the accommodation described above is very different from that which arises in two-stage “accommodation” games (see pp. 328–329 of Tirole 1988), in which firms take some action in a first stage (e.g., restrict capacity) so as to soften price competition in a second stage. While the



**Fig. 6** Transient distributions  $\mu_t^m(\omega^m)$  for  $t = 30$  and  $t = 100$  for  $\sigma = 0.1$  (left panel) and  $\sigma = 2$  (right panel)

accommodation in such two-stage games is predicated only on irreversible investments made in the past, phases of accommodation depend critically on firms' expectations for the future; specifically, firms' accommodate one another because they both benefit from slowing the rate at which the industry evolves toward an asymmetric industry structure.

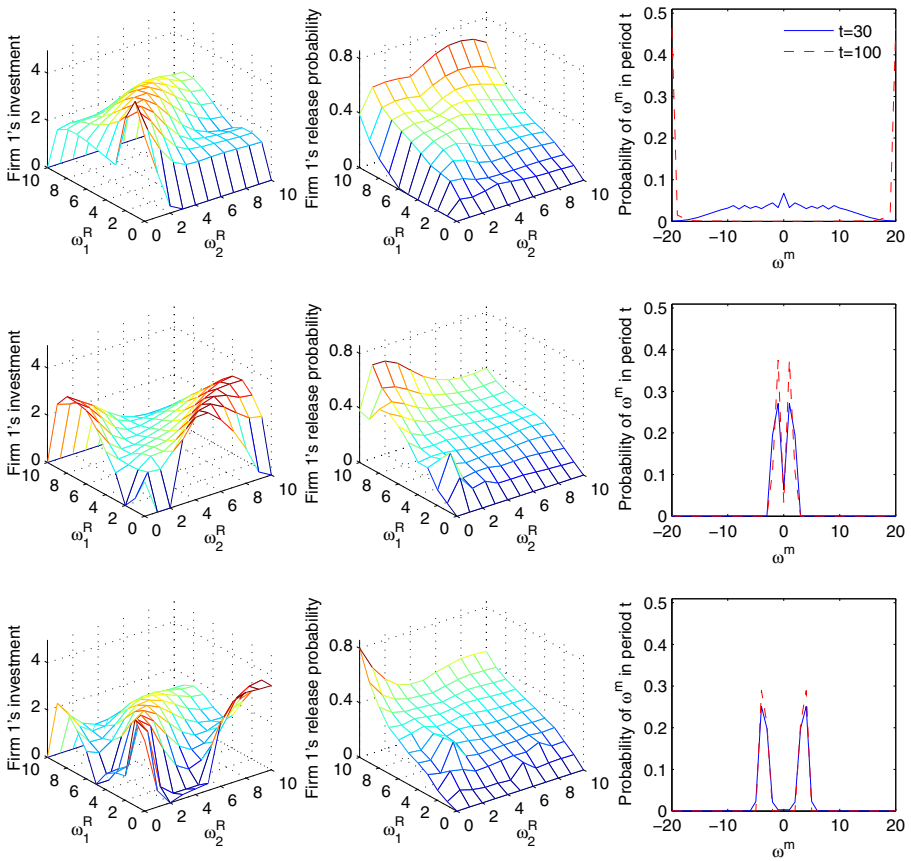
### 6.1.3 A high degree of horizontal differentiation

To explore the effect of a high degree of horizontal differentiation, I examine the three equilibria that I have found for the  $\sigma = 4$  parameterization. Each row of Fig. 7 presents one equilibrium; the  $\omega^m = 0$  cross-sections of the policy functions are presented in the left and middle columns, and the corresponding transient distributions for periods 30 and 100 are presented in the right column.<sup>34</sup> I briefly describe the equilibrium in the top row before moving on to a discussion of the *asymmetric R&D wars* that characterize the equilibria in the middle and bottom rows.

The equilibrium in the top row is qualitatively similar to the accommodative equilibrium presented in Fig. 5. While the  $\omega^m = 0$  cross-section includes only one accommodative trench (at  $\omega_2^R = 0$ ), other cross-sections include additional accommodative trenches. Like the accommodative equilibrium presented above, this equilibrium ultimately yields an extremely asymmetric industry structure. The size of the leader's expected lead after 30 (100) periods is 6.82 (19.76).

**Asymmetric R&D wars** The equilibria presented in the middle and bottom rows of Fig. 7 are characterized by asymmetric R&D wars. I first describe these equilibria and then explain why asymmetric R&D wars arise.

<sup>34</sup>I plot the  $\omega^m = 0$  cross-sections of only firm 1's policy functions; as in Figs. 3 and 5, the cross-sections of firm 2's policy functions are simply the transposes of those presented for firm 1.



**Fig. 7** Three equilibria ( $\sigma = 4$ ); one in each row.  $\omega^m = 0$  cross-sections of policy functions for R&D investment  $x_1(\omega)$  (left panels) and release probability  $r_1(\omega)$  (middle panels). Transient distributions  $\mu_t^m(\omega^m)$  for  $t = 30$  and  $t = 100$  (right panels)

Unlike the equilibria of the full model presented thus far, the equilibrium in the middle row of Fig. 7 admits no preemption whatsoever; i.e., on or near the diagonal—where neither firm has too large an R&D stock advantage—the R&D investment levels and updating probabilities of both firms are relatively low. Firms do however engage in *asymmetric R&D wars*; once one firm achieves a sufficiently large R&D stock advantage, both firms drastically increase their R&D investments. Hence, while firms do not engage in aggressive competition *for* the lead, they do engage in aggressive competition *once* one firm achieves a large lead. Because the follower invests more than the leader,<sup>35</sup> it tends to succeed in narrowing the leader’s lead and, as a

<sup>35</sup>The panel in the middle row and middle column of Fig. 7 shows that when a firm gains a large R&D stock lead, it updates with high probability. Hence, a follower invests heavily in R&D in hopes of preventing the leader from gaining such a large lead. Because the follower updates with low probability, the leader faces no similar incentive.

result, a (nearly) symmetric industry structure arises; the size of the leader's expected lead after 30 (100) periods is 1.36 (1.27). The aggressive R&D investment competition that occurs when one firm gains a sufficiently large lead both induces firms to avoid striving for a large lead and prevents a firm from sustaining such a lead if it is achieved.<sup>36</sup>

The equilibrium in the bottom row of Fig. 7 is qualitatively similar to the one in the middle row, with one exception; firms engage in preemption by investing heavily in R&D when neither firm has too large a lead. The preemptive investment behavior gives rise to an industry structure that is more asymmetric than that of the equilibrium in the middle row; the size of the leader's expected lead after 30 (100) periods is 3.55 (3.60). However, the industry structure is still relatively symmetric because, as in the equilibrium in the middle row, asymmetric R&D wars prevent either firm from gaining or sustaining a large lead.<sup>37</sup>

**Why do firms engage in asymmetric R&D wars?** I next explain why asymmetric R&D wars arise. In the preemptive equilibrium (Section 6.1.1) and the accommodative equilibrium (Section 6.1.2), a firm that falls behind accepts its role as the laggard, i.e., it invests in product development only in order to slow its decline in the product market. However, in an asymmetric R&D war, a laggard fights back by investing heavily in R&D in hopes of narrowing the leader's lead (and this gives rise to aggressive competition because the leader responds by also investing heavily). As explained in Section 4, an increase in the degree of horizontal differentiation weakens the leader's (short-run) incentives to invest in product development, and strengthens those of the follower. When  $\sigma = 4$ , the leader's (short-run) incentives are sufficiently weak and the follower's are sufficiently strong such that it is worthwhile for the follower to invest in catching up. However, R&D stockpiling and endogenous version releases also play a critical role. In fact, Section 5 shows that an otherwise identical model *without* R&D stockpiling and endogenous version releases does *not* admit asymmetric R&D wars.<sup>38</sup> This raises the question of why asymmetric R&D wars arise only when firms can stockpile R&D and decide when to release new versions.

In the full model, a firm that falls behind in the R&D stock race does not necessarily fall behind in the product market; rather, it falls behind in the product market only if its rival releases a new version. However, in the benchmark model, a firm that falls behind immediately becomes the product market laggard, and accordingly

<sup>36</sup>In contrast, in earlier quality ladder models (e.g., Borkovsky et al. 2012) and more generally in dynamic models of industry evolution (e.g., Besanko et al. 2010), symmetric industry structure tends to arise when firms compete *less* aggressively.

<sup>37</sup>The equilibria in the middle and lower panels of Fig. 7 are also characterized by accommodation. While the  $\omega^m = 0$  cross sections in the lower panels do not include any accommodative trenches and the ones in the middle panels include only one (at  $\omega_2^R = 0$  for  $\omega_1^R \geq 9$ ), there are several prominent accommodative trenches in other cross-sections for both equilibria.

<sup>38</sup>The bottom panels of Fig. 2 show that in the benchmark model, when the degree of horizontal differentiation is high, as the laggard falls behind it simply reduces its R&D investment and the leader increases its R&D investment (for  $\omega^m \leq 10$ ), making it extremely unlikely that the laggard narrows the leader's lead; accordingly, an asymmetric industry structure arises.

it necessarily finds itself on the flatter portion of the (static) profit function. It follows that falling behind (in the R&D stock race) in the full model does not weaken investment incentives as much as falling behind (in the product market) in the benchmark model. To put it more simply, in the full model, it is as if the R&D stock laggard strives to catch up to the leader *before* the leader releases a new version in order to avoid falling behind in the product market; in the benchmark model, the laggard has no such opportunity.


Finally, because the model admits qualitatively different equilibria at the  $\sigma = 4$  parameterization (and more generally when the degree of horizontal differentiation is sufficiently high, see Fig. 10 in the Appendix), it follows that the nature of firms' strategies and accordingly the industry structure that arises are determined not only by the parameterization, but also by the equilibrium itself.<sup>39</sup>

## 6.2 Increasing the cost of releasing a new version

The cost of releasing a new version could increase for a wide variety of reasons, e.g., an increase in the complexity of incorporating product improvements into a new version, an increase in the cost of labor or capital that must be employed to launch a new version, or an increase in the marketing costs that must be incurred. I explore how an increase in the cost of releasing a new version affects firms' product development strategies. Because the equilibria presented in Sections 6.1.1 and 6.1.2 are characterized by the two most prevalent behaviors that arise—preemption races and phases of accommodation—I use them as baselines. That is, I explore how these two behaviors change as the mean release cost increases. To this end, I use each equilibrium as a starting point for a simple continuation method in which I incrementally increase the mean release cost  $E[\phi_i] = (G_l + G_u)/2$  by increasing both  $G_l$  and  $G_u$ , while holding the width of the support,  $G_u - G_l$ , constant. In Fig. 8, I summarize each equilibrium computed using the size of the leader's expected lead in period 30,  $L^{30}$ .

I begin by exploring how the accommodative equilibrium presented in Section 6.1.2 changes as the mean release cost increases. In the left panel of Fig. 8,  $L^{30}$  increases non-monotonically in the mean release cost. This non-monotonic relationship arises because increasing the release cost introduces two countervailing effects. The *direct effect* is intuitive; an increase in the release cost reduces the returns to innovation and accordingly weakens firms' incentives to innovate. Therefore, firms invest less in R&D and release new versions less frequently. It follows that a market leader marches out toward industry dominance more slowly and accordingly the size of the leader's expected lead in period 30 decreases. If this were the only effect at play, the size of the leader's expected lead would be strictly decreasing in the mean release cost.

<sup>39</sup>Besanko et al. (2010) explain why their model of dynamic price competition admits multiple equilibria and, specifically, how different equilibria are underpinned by different beliefs about the future (see pp. 492–493). The same explanation applies to my model.



**Fig. 8** Leader's expected lead at  $t = 30$  vs. the mean release cost for  $\sigma = 2$  (left panel) and  $\sigma = 0.1$  (right panel)

The *strategic effect* is perhaps less intuitive and arises strictly because of oligopolistic competition. As the release cost increases, it becomes more costly for a product market follower to catch up to the leader. This weakens the follower's incentive to innovate. As a result, the higher the release cost, the less likely it is that the follower ever catches up to the leader. So, a higher release cost makes the position of a product market leader more secure. This induces firms to compete more aggressively for the lead by engaging in more aggressive preemption. That is, in the heat of a preemption race, a firm sometimes updates with higher probability—despite the higher cost—and invests more heavily in R&D. At the same time, phases of accommodation become less prevalent. For a sufficiently low mean release cost, the strategic effect dominates the direct effect. As a result, as the release cost increases, the industry converges to an asymmetric industry structure more quickly. The left panel of Fig. 8 shows that this occurs for  $E[\phi_i] \leq 66.5$ . However, for a sufficiently high release cost ( $E[\phi_i] > 66.5$ ), the direct effect dominates and accordingly the expected lead in period 30 declines as the release cost increases.<sup>40</sup>

I next explore how the preemptive equilibrium presented in Section 6.1.1 changes as the mean release cost increases. The right panel of Fig. 8 shows that while the leader's expected lead is non-monotonic in the release cost, the non-monotonicity is much milder than the one in the left panel. As the mean release cost increases, the direct effect quickly overwhelms the strategic effect and causes the size of the leader's expected lead to decrease. This is because the preemptive equilibrium that serves as the starting point is already characterized by intense preemption and does not admit phases of accommodation. Therefore, the strategic effect—which intensifies preemption and lessens the extent of accommodation—has only limited impact before it is overwhelmed by the direct effect.

<sup>40</sup>The relatively small non-monotonicities in the left panel of Fig. 8 arise because of slight qualitative changes in the equilibrium policy functions.

## 7 Discussion & conclusion

The purpose of this paper is to better understand product development that is characterized by repeated releases of new versions of existing products. To this end, this paper presents the first oligopolistic model of product development that endogenizes firms' decisions over how much to invest in R&D to develop product enhancements and when to release new versions that incorporate them. I use the model to explore the role of two key industry fundamentals—the degree of horizontal product differentiation and the cost of releasing a new version—on firms' product development strategies and accordingly the evolution of industry structure.

I find that varying the degree of horizontal differentiation gives rise to three distinctly different types of competitive dynamics: *preemption races* when the degree of horizontal differentiation is low; *phases of accommodation* when it is moderate; and *asymmetric R&D wars* when it is high. Furthermore, I show that phases of accommodation and asymmetric R&D wars arise *only* in a model that incorporates R&D stockpiling and (endogenous) version releases. My results also emphasize the strategic nature of R&D stockpiling, i.e., firms stockpile R&D not only because it can ultimately enhance profits, but also because it can be used to influence rival behavior.

I also explore the effect of the cost of releasing a new version on firms' product development strategies. I find that an increase in the release cost can induce firms to compete *more* aggressively for the lead and, in doing so, release new versions *more* frequently despite the higher cost. This result yields a useful managerial insight. There are various reasons why an exogenous change in the cost of releasing new versions might occur. For example, in recent years, China has made efforts to develop a more high-tech economy and, as a result, some western firms have begun to build research labs there (Bradsher 2010; Hout and Ghemawat 2010). If firms anticipate that they will be able to outsource some of their R&D to China, this may ultimately reduce the cost of releasing a new version. One might be inclined to think that a decrease in the release cost would induce firms to release new versions more frequently and therefore intensify competition. However, the insights discussed in Section 6.2 explain why the exact opposite might occur; i.e., a decrease in the release cost could soften competition to innovate because it would make any product market lead less secure.

While the model presented in this paper perhaps best applies to R&D-intensive CPG categories, it could be augmented to study other types of categories. Consider the Internet browser industry. From 1995 to 2001, Microsoft and Netscape engaged in the first so-called *browser war* by both investing heavily to develop new features and releasing new versions of their respective browsers, which incorporated these features, on a frequent basis. Collectively, they released twelve versions in a span of only six years. By 2001, Microsoft's Internet Explorer (IE) had won the browser war and come to dominate the market with a usage share of 90% (Markoff 2001). This brought an end to both the rapid innovation and the frequent version releases (Wildstrom 2003); Microsoft would not release another version of its browser for five years. In fact, in 2003, Microsoft announced that it would cease to release standalone versions of IE and that future enhancements would be bundled with operating system upgrades (Hansen 2003). However, in the face of increasing competition with

the Mozilla Firefox browser—which included new features not offered by IE and had begun to encroach on IE’s usage share—Microsoft modified its strategy by releasing a new standalone version, IE7, in 2006. This version introduced several features that were already offered by Firefox and seemed to close the gap in perceived quality that had existed between IE and Firefox (Hoover 2006). Moreover, Microsoft announced its plan to release the next version of IE within 18 months (Hoover 2006), and ultimately released it after 29 months (Fried 2009). Therefore, in the face of the threat posed by Firefox, Microsoft began to release versions with greater frequency. While competition in this industry has been complicated by a host of other issues—such as the *United States v. Microsoft* antitrust case and the issues explored therein—strategically timed version releases that incorporate improvements developed through R&D have played a prominent role. Most notably, as explained above, Microsoft and Netscape engaged in a preemption race that was characterized by intense investment and frequent version releases until Microsoft came to dominate the browser market. This preemption race is similar in nature to those described in Section 6.1.1. However, while those preemption races can be attributed to a low degree of horizontal differentiation, the preemption race in the browser category likely arose because of the presence of indirect network effects.<sup>41</sup> Therefore, to apply this model to the browser category, it would be important to formally incorporate indirect network effects.<sup>42</sup>

Another R&D-intensive CPG category that is characterized by periodic releases of new and improved versions is men’s razors with replacement blades. The market-leading brands, Gillette and Schick, have invested heavily in R&D to increase the number of blades, develop new ways of pivoting the blade head, increase lubrication, reduce friction, and improve razor grip (Richardson 2010; Chain Drug Review 2010). This category however also possesses two other important characteristics (Hartmann and Nair 2010): (i) razors are durable; and (ii) razors and blades are “tied”, i.e., a razor can only be used with compatible blades produced by the same manufacturer. Therefore, to model product development in this category, one would allow consumers to make forward-looking razor adoption and replacement decisions (Gordon 2009; Goettler and Gordon 2011). Furthermore, one would allow firms to make forward-looking product development and pricing decisions for both razors and blades, taking into consideration the implications of product complementarity for the demand of each product.<sup>43</sup>

As with any theoretical work, this paper has limitations. First, as illustrated by the above examples, some R&D-intensive consumer product categories are complicated by other important characteristics such as indirect network effects, product durability, and the tied nature of products. If one is interested in studying one such

---

<sup>41</sup>Indirect network effects existed because web content developers preferred to produce content for browsers with large user bases and users preferred browsers for which much compatible content was available (Bresnahan 2001).

<sup>42</sup>Nair et al. (2004), Markovich (2008), Markovich and Moenius (2009), Dubé et al. (2010), and Lee (2013) explore the impact of indirect network effects in an oligopolistic context.

<sup>43</sup>To incorporate product durability into the model, in addition to including forward-looking consumers and dynamic pricing, one would augment the state space to include the ownership distribution across products (Goettler and Gordon 2011).



category in particular, then it might be important to formally account for the relevant characteristic.

Second, in some categories—e.g., razors and blades—a manufacturer continues to sell the old version of a product even after the new version is released. I have abstracted from this because my objective is to study firms' product development strategies; once a new version is released, even if a firm continues to sell the old version, it rarely invests in improving it. Moreover, while it would be theoretically straightforward to incorporate this into the above model, it would significantly increase the computational burden because it would add several dimensions to the state space (one for each old product that continues to be sold).

Third, the model in this paper does not include entry and exit. If one is interested in studying competition between a small number of national brands that dominate a market that has seen little if any entry or exit—like the U.S. diaper market (discussed in Section 1)—then one might abstract from entry and exit. Otherwise, one might incorporate entry and exit in order to explore how firms strive to prevent entry and induce exit via their product development strategies.<sup>44</sup> (Borkovsky et al. (2012) show that in the presence of entry and exit, firms investing in R&D to enhance product quality engage in both *limit investment* and *predatory investment*.)

Fourth, some categories are characterized by both repeated releases of new versions of existing products *and* product line extensions. In such categories, a firm might test an innovation by introducing it via a product line extension.<sup>45</sup> Only if it proves to be successful would the firm later incorporate the innovation into existing products.<sup>46</sup> By abstracting away from product line extensions, I have been able to devise a tractable model of product development via new releases of existing products. Incorporating both new releases of existing products *and* product line extensions into one model would raise a number of formidable challenges, one being that the model would have to allow for multi-product firms. This would significantly increase the dimensionality of the state space, and accordingly the computational burden of the model, just as in the other multi-product firm scenario (tied products) discussed above.

Finally, my dynamic model incorporates a static model of product market competition that gives rise to a monotonically increasing profit function (see Fig. 1). This

---

<sup>44</sup>Incorporating entry and exit into the model would be theoretically straightforward. However, it would increase computational burden because one would not be able to restrict attention to the difference between firms' respective product qualities. If one firm exits, the remaining incumbent firm would have to be characterized by its product's absolute quality. Therefore, it would be necessary to retain absolute as opposed to relative product qualities in the state space.

<sup>45</sup>The notion of learning about uncertain demand from observed sales is explored in Hitsch (2006).

<sup>46</sup>In the diaper category, several innovations have been introduced via product line extensions, e.g., gel technology (Pampers 1986), Velcro tabs (Huggies 1993, Pampers 1994), and stretchable side panels (Pampers 1994) (Parry and Jones 2001). These innovations were later broadly incorporated into the firms' respective product lines. Both P&G and Kimberly-Clark have also introduced many innovations directly into new versions of existing products, e.g., tape closures (Pampers 1971), stay-dry lining (Pampers 1976), a waist shield (Pampers 1985), and contoured elastic leg bands (Huggies 1992) (Dyer et al. 2004; Parry and Jones 2001).

is the standard approach in the literature on quality ladder models.<sup>47</sup> It follows that in terms of static profits, a firm only benefits from climbing up the quality ladder and only suffers from falling down it. However, in the literature on endogenous quality choice<sup>48</sup>—where firms typically choose product quality in a first stage and then engage in price competition in a second stage (e.g., Shaked and Sutton 1982, Moorthy 1988, Neven and Thisse 1990)—a firm's profit function can be non-monotonic in the size of its quality advantage. This stems from the assumption that consumers are heterogeneous in terms of marginal utility of quality. (In my model, consumers differ only in terms of their additive shocks.) Therefore, vertical product differentiation softens price competition—because the high (low) quality firm serves the consumers with higher (lower) taste for quality—and increases the profits of *both* firms. Incorporating consumer heterogeneity in terms of taste for quality into my model—and more broadly into the literature on quality ladder models—has the potential to yield very interesting results; because a firm obtains some benefit from falling behind, it would face different incentives when making its product development decisions. I feel that this is a fertile area for future research.<sup>49</sup>

Notwithstanding these limitations, this paper helps us understand product development in categories characterized by repeated releases. By devising a model of such product development in the Ericson and Pakes (1995) framework and solving it using numerical methods, I have been able to explore an important type of product development that is not accessible through more familiar analytic research methods. The equilibrium behaviors that arise in this model are qualitatively different from those that arise in earlier papers on product development and R&D competition, which do not endogenize firms' version release decisions. This shows that it is important not to abstract away from this aspect of competition when studying categories in which it is prevalent.

**Acknowledgments** I am extremely grateful to Mark Satterthwaite, David Besanko, Ulrich Doraszelski, Sarit Markovich, Andrew Ching and Avi Goldfarb for their guidance and feedback. I also thank Guy Arie, Allan Collard-Wexler, J.P. Dubé, Ron Goettler, Paul Grieco, Delaine Hampton, Dan Horsky, Yaroslav Kryukov, Mitch Lovett, Lauren Lu, Karl Schmedders, audiences at numerous conferences and universities, and especially the editor and reviewers for very helpful comments. I gratefully acknowledge financial support from the General Motors Center for Strategy in Management at Northwestern's Kellogg School of Management and SSHRC Grant 410-2011-0356. I thank Marina Milenkovic for superb research assistance.

<sup>47</sup>The literature on analytically tractable quality ladder models (see footnote 5) assumes that an exogenous profit function is (weakly) increasing in the size of a firm's lead. In the empirical (Gowrisankaran and Town 1997; Goettler and Gordon 2011; Borkovsky et al. 2017) and computational applied theory (Pakes and McGuire 1994; Markovich 2008; Markovich and Moenius 2009; Borkovsky et al. 2012; Goettler and Gordon 2014) literatures, the (endogenous) static profit function is strictly increasing in a firm's quality because—as in my model—the marginal utility of quality is constant across consumers.

<sup>48</sup>The literature on endogenous quality choice explores the decisions that firms face in positioning new products relative to their competitors (see p. 141 of Moorthy 1988). The literature on quality ladder models explores competition amongst firms that repeatedly invest in R&D to incrementally enhance (or sustain) the quality of existing products while also repeatedly engaging in product market competition.

<sup>49</sup>This insight stems from Referee 2's comments, for which I am extremely grateful.

## Appendix

In Sections A.1 and A.2, I present details on the dynamic model and on computation that have been omitted from the paper. In Section A.3, I present the equilibrium of the static product market game, including the equilibrium market share functions that have been omitted from Fig. 1 in the paper. In Section A.4, I present a detailed version of the benchmark model that is described in Section 5 of the paper. In Section A.5, I summarize all of the equilibria that I have computed for a fine discretization of a wide range of  $\sigma$  (i.e., degree of horizontal differentiation) values. In Section A.6, I present an alternative version of the static model of product market competition that is based on the Hotelling (1929) “linear city” model. I discuss the benefits of using the Logit model presented in Section 2, as opposed to the Hotelling model, in the baseline specification. I then show that a version of the dynamic model that includes the Hotelling model as its period game admits equilibria that are qualitatively similar to those of the baseline model in Section 2. In Section A.7, I present equilibria of two alternative versions of the dynamic model; each alternative version “turns off” one key feature of the dynamic model so as to help readers better understand its role. In Section A.8, I discuss an example that shows how introducing uncertainty into the version release process mitigates the effects of the edge of the state space.

### A.1 The model: deriving the optimality conditions

In this section, I present the optimization problems that firms face and derive the optimality conditions. An abridged version of this material appears in Sections 2.1 and 2.3.

**State-to-state transitions** In each period, the firms’ respective updating and R&D investment decisions determine the industry state that arises in the next period. As explained above, a firm is able to enhance the quality of its product by releasing a new version, which incorporates each unit of the firm’s R&D stock with some probability. Specifically, if firm  $i$  possesses  $\omega_i^R$  units of R&D stock and it releases a new version, its product quality improves by  $\bar{\omega}_i^R \in \{0, 1, \dots, \omega_i^R\}$  units with probability  $s(\bar{\omega}_i^R | \omega_i^R)$ . This resets firm  $i$ ’s R&D stock,  $\omega_i^R$ , to zero. It follows that:

1. if firm 1 releases a new version and firm 2 does not, the industry state transitions from  $(\omega^m, \omega_1^R, \omega_2^R)$  to  $(\min(\omega^m + \bar{\omega}_1^R, L^m), 0, \omega_2^R)$  with probability  $s(\bar{\omega}_1^R | \omega_1^R)$ ;
2. if firm 2 releases a new version and firm 1 does not, the industry state transitions from  $(\omega^m, \omega_1^R, \omega_2^R)$  to  $(\max(\omega^m - \bar{\omega}_2^R, -L^m), \omega_1^R, 0)$  with probability  $s(\bar{\omega}_2^R | \omega_2^R)$ ; and
3. if each firm releases a new version, the industry state transitions from  $(\omega^m, \omega_1^R, \omega_2^R)$  to  $(\min(\max(\omega^m + \bar{\omega}_1^R - \bar{\omega}_2^R, -L^m), L^m), 0, 0)$  with probability  $s(\bar{\omega}_1^R | \omega_1^R) \times s(\bar{\omega}_2^R | \omega_2^R)$ .

The max and min functions are included above simply to ensure that transitions are confined to the state space.

In subperiod 2, the industry is initially in state  $\omega'$ . Firm  $i$ 's R&D stock for the subsequent period,  $\omega_i^{R''}$ , is determined by the stochastic outcome of its investment decision:

$$\omega_i^{R''} = \omega_i^{R'} + v_i,$$

where  $v_i \in \{0, 1\}$  is a random variable governed by firm  $i$ 's investment  $x_i \geq 0$ . If  $v_i = 1$ , the investment is successful and the firm  $i$ 's R&D stock increases by one. The probability of success is  $\frac{\alpha x_i}{1 + \alpha x_i}$ , where  $\alpha > 0$  is a measure of the effectiveness of investment. To simplify exposition, I define

$$q(v_i | x_i, \omega_i^R) = \begin{cases} \frac{\alpha x_i}{1 + \alpha x_i} & \text{if } v_i = 1, \\ \frac{1}{1 + \alpha x_i} & \text{if } v_i = 0, \end{cases}$$

if  $\omega_i^R \in \{0, 1, \dots, L^R - 1\}$ , and  $q(v_i = 0 | x_i, L^R) = 1$ , which simply enforces the bounds of the state space.

**Bellman equation** To derive the Bellman equation, I first consider the investment decisions that firms make in subperiod 2 and then the updating decisions that they make in subperiod 1. I let  $V_i(\omega, \phi_i)$  denote the expected net present value of all future cash flows to firm  $i$  in state  $\omega$  in subperiod 1, immediately after it has drawn release cost  $\phi_i$ . Because I solve for a symmetric equilibrium (as explained in Section 2.3), I restrict attention to firm 1's problem.

**Investment decision** At the beginning of subperiod 2, the industry is in state  $\omega'$ . The expected net present value of cash flows to firm 1 is

$$U_1(\omega') \equiv \max_{x_1 \geq 0} \left\{ -x_1 + \beta E \left[ V_1(\omega'', \phi'_1) | \omega', x_1 \right] \right\}. \tag{2}$$

The continuation value is

$$E[V_1(\omega'', \phi'_1) | \omega', x_1] = \sum_{v_1} W_1(v_1 | \omega') q(v_1 | x_1, \omega_1^{R'}),$$

where

$$W_1(v_1 | \omega') \equiv \sum_{v_2 \in \{0,1\}} q(v_2 | x_2(\omega'), \omega_2^{R'}) \int_{\phi'_1} V_1 \left( (\omega^m, \omega_1^R + v_1, \omega_2^{R''}), \phi'_1 \right) dG(\phi'_1)$$

is the expectation of firm 1's value conditional on an investment success ( $v_1 = 1$ ) or failure ( $v_1 = 0$ ), and  $x_2(\omega')$  is firm 2's R&D investment in industry state  $\omega'$ . Firm 1 chooses investment  $x_1 \geq 0$  that maximizes the expected net present value of its future cash flows. Solving firm 1's optimization problem on the right-hand side of Eq. 2, I find that

$$x_1(\omega') = \max \left\{ 0, \frac{-1 + \sqrt{\beta \alpha [W_1(1 | \omega') - W_1(0 | \omega')]} }{\alpha} \right\} \tag{3}$$

if  $W_1(1 | \omega') \geq W_1(0 | \omega')$ , and  $x_1(\omega') = 0$  otherwise, for all  $\omega'$  such that  $\omega_1^R \neq L^R$ ; and  $x(\omega') = 0$  for all  $\omega'$  such that  $\omega_1^R = L^R$ .

**Version release decision** Having solved for firm 1’s optimal investment in subperiod 2, I can now solve for firm 1’s optimal updating decision in subperiod 1. Consider an industry that is in state  $\omega$  in subperiod 1, after firms have drawn their respective costs of updating. Let  $Y_1^1(\omega)$  be the expected net present value of all future cash flows to firm 1 in state  $\omega$  if it decides to update and firm 2 does not; the period profits that firm 1 earns in state  $\omega$  and the release cost it incurs are not incorporated into  $Y_1^1(\omega)$ , but are included below in Bellman Eq. 7. Define  $Y_1^2(\omega)$  similarly, and let  $Y_1^{12}(\omega)$  be defined analogously for the case in which both firms choose to update. It follows that

$$Y_1^1(\omega^m, \omega_1^R, \omega_2^R) \equiv \sum_{i=0}^{\omega_1^R} s(i|\omega_1^R) U_1(\min(\omega^m + i, L^m), 0, \omega_2^R), \tag{4}$$

$$Y_1^2(\omega^m, \omega_1^R, \omega_2^R) \equiv \sum_{j=0}^{\omega_2^R} s(j|\omega_2^R) U_1(\max(\omega^m - j, -L^m), \omega_1^R, 0), \tag{5}$$

$$Y_1^{12}(\omega^m, \omega_1^R, \omega_2^R) \equiv \sum_{i=0}^{\omega_1^R} \sum_{j=0}^{\omega_2^R} s(i|\omega_1^R) s(j|\omega_2^R) \times U_1(\min(\max(\omega^m + i - j, -L^m), L^m), 0, 0). \tag{6}$$

Let firm 1’s perceived probability that firm 2 releases a new version of its product, conditional on the current industry state  $\omega$  be  $r_2(\omega)$ . Firm 1’s value function  $V_1 : \Omega \times \Phi \rightarrow \mathbb{R}$  is implicitly defined by the Bellman equation

$$V_1(\omega, \phi_1) = \pi_1(\omega^m) + \max \left\{ (1 - r_2(\omega))U_1(\omega) + r_2(\omega)Y_1^2(\omega), \right. \\ \left. -\phi_1 + (1 - r_2(\omega))Y_1^1(\omega) + r_2(\omega)Y_1^{12}(\omega) \right\}. \tag{7}$$

Substituting  $x_1(\omega')$  from Eq. 3 into Eq. 2;  $U_1(\omega')$  from Eq. 2 into Eqs. 4–6; and  $Y_1^1(\omega)$ ,  $Y_1^2(\omega)$ , and  $Y_1^{12}(\omega)$  from Eqs. 4–6 into Bellman Eq. 7, I determine whether firm 1 chooses to update the good being sold in the product market in industry state  $\omega$  when it draws release cost  $\phi_1$ . Letting  $\chi_1(\omega, \phi_1)$  be the indicator function that takes the value of one if firm 1 chooses to update, and zero otherwise, yields

$$\chi_1(\omega, \phi_1) = \arg \max_{\chi \in \{0,1\}} (1 - \chi) \left\{ (1 - r_2(\omega))U_1(\omega) + r_2(\omega)Y_1^2(\omega) \right\} \\ + \chi \left\{ -\phi_1 + (1 - r_2(\omega))Y_1^1(\omega) + r_2(\omega)Y_1^{12}(\omega) \right\}.$$

The probability of drawing a  $\phi_1$  such that  $\chi_1(\omega, \phi_1) = 1$  determines the probability of updating or

$$r_1(\omega) = G \left( \left\{ (1 - r_2(\omega))Y_1^1(\omega) + r_2(\omega)Y_1^{12}(\omega) \right\} \right. \\ \left. - \left\{ (1 - r_2(\omega))U_1(\omega) + r_2(\omega)Y_1^2(\omega) \right\} \right). \tag{8}$$

I can now restate the Bellman equation in terms of firm 1’s expected value function

$$V_1(\omega) \equiv \int_{\phi_1} V_1(\omega, \phi_1) dG(\phi_1).$$

The expected value function  $V_1 : \Omega \rightarrow \mathbb{R}$  is implicitly defined by the Bellman equation

$$V_1(\omega) = \pi_1(\omega^m) + (1 - r_1(\omega)) \left\{ (1 - r_2(\omega))U_1(\omega) + r_2(\omega)Y_1^2(\omega) \right\} + r_1(\omega) \left\{ -\phi_1(\omega) + (1 - r_2(\omega))Y_1^1(\omega) + r_2(\omega)Y_1^{12}(\omega) \right\}, \tag{9}$$

where  $\phi_1(\omega)$  is the expectation of  $\phi_1$  conditional on updating in state  $\omega$ .

**Equilibrium** I restrict attention to symmetric Markov perfect equilibria in pure strategies. Theorem 1 in Doraszelski and Satterthwaite (2010) establishes that such an equilibrium always exists.<sup>50</sup> In a symmetric equilibrium, the investment decision taken by firm 2 in state  $\omega$  is identical to the investment decision taken by firm 1 in state  $\omega^{[2]} \equiv (-\omega^m, \omega_2^R, \omega_1^R)$ , i.e.,  $x_2(\omega) = x_1(\omega^{[2]})$ . A similar relationship holds for the probability of releasing a new version and the value function:  $r_2(\omega) = r_1(\omega^{[2]})$  and  $V_2(\omega) = V_1(\omega^{[2]})$ .<sup>51</sup> It therefore suffices to determine the value and policy functions of firm 1. Solving for an equilibrium for a particular parameterization of the model amounts to finding a value function  $V_1(\cdot)$  and policy functions  $x_1(\cdot)$  and  $r_1(\cdot)$  that satisfy the Bellman Eq. 9 and the optimality conditions 3 and 8.

### A.2 Computation

This section complements Section 3. I present the formal specification of the distribution that determines how many units of R&D stock are successfully incorporated into a new version. I also present additional detail on the algorithm used to compute equilibria.

**Uncertainty in the version release process** Recall that if a firm with R&D stock  $\omega_i^R$  updates, its product quality improves by  $\bar{\omega}_i^R \in \{0, 1, \dots, \omega_i^R\}$  with probability  $s(\bar{\omega}_i^R | \omega_i^R)$ . I assume that  $s(\cdot | \omega_i^R)$  is a generalized version of the binomial distribution; in particular,  $s(\bar{\omega}_i^R | \omega_i^R)$  is the probability of obtaining exactly  $\bar{\omega}_i^R$  successes out of  $\omega_i^R$  Bernoulli trials that are independent but are *not* identically distributed. The success

<sup>50</sup>The version release decision in this model is analogous to the entry and exit decisions in Doraszelski and Satterthwaite (2010); by assuming that the release cost is random and privately known, I “purify” the mixed strategy equilibria that the model would admit were the release cost fixed (Harsanyi 1973).

<sup>51</sup>Because I restrict attention to symmetric equilibria, all relevant differences between firms are reflected in the industry state  $\omega$ . Accordingly each firm’s behavior is a function of only the size of its product market advantage/disadvantage and the status of the R&D stock race. For example, firm 1’s behavior in industry state (7,4,2) is identical to firm 2’s behavior in industry state (-7,2,4); in each case, the firm in question has a product market lead of 7 and has an R&D stock lead of 4 to 2.

probabilities of the  $\omega_i^R$  Bernoulli trials are  $q_1, q_2, \dots, q_{\omega_i^R}$ , respectively. Let  $\mathbf{q}_{\omega_i^R} \equiv (q_1, q_2, \dots, q_{\omega_i^R})$  for  $\omega_i^R = 1, \dots, L^R$ , and let

$$q_n = \rho^n$$

for  $n = 1, \dots, \omega_i^R$ , where  $\rho \in (0, 1)$ . So, the probability that firm  $i$  succeeds in incorporating the first unit of its R&D stock into a new version of its product is  $\rho$ , the probability that it succeeds in incorporating the second unit is  $\rho^2 < \rho$  etc. That is, each additional unit of R&D stock is less likely to be successfully incorporated into a new version than the previous unit.<sup>52</sup>

**Algorithm** I compute equilibria using the Gauss-Jacobi version of the Pakes and McGuire (1994) algorithm. To explore the equilibrium correspondence, I nest the Pakes and McGuire (1994) algorithm in a simple continuation method (Judd 1998). The simple continuation method computes one equilibrium for each of a sequence of parameterizations of the model. For example, to compute equilibria for  $\sigma \in \{\sigma_1, \sigma_2, \dots, \sigma_T\}$  where  $\sigma_{t-1} \leq \sigma_t$  and  $\sigma_{t-1} \approx \sigma_t$ , while holding all other parameters fixed, one first solves for an equilibrium for  $\sigma = \sigma_1$ . This equilibrium serves as the starting point for the simple continuation method. The simple continuation method sequentially solves for an equilibrium for each of  $\sigma_2, \dots, \sigma_T$ . For each  $\sigma_t$ , it uses the equilibrium computed for  $\sigma_{t-1}$  as the starting point for the Pakes and McGuire (1994) algorithm. In this sense, the simple continuation method provides a systematic approach to selecting starting points for the Pakes and McGuire (1994) algorithm. Running the simple continuation method upward—from  $\sigma_1$  to  $\sigma_T$ —and then downward—from  $\sigma_T$  to  $\sigma_1$ —can help identify multiple equilibria; if the model admits multiple equilibria for a given  $\sigma$  value, then the simple continuation method might compute a different equilibrium when approaching  $\sigma$  from the left than it does when approaching  $\sigma$  from the right.

### A.3 Results: product market competition

The (static) equilibrium price and profit functions for low ( $\sigma = 0.1$ ), moderate ( $\sigma = 2$ ), and high ( $\sigma = 4$ ) degrees of horizontal differentiation are presented in Fig. 1 in the paper. Because Fig. 1 does not include the equilibrium market share functions—which play a central role in the discussion in Section 4—I include them in Fig. 9.

<sup>52</sup>For this specification, I have found that setting  $\rho = 0.95$  mitigates the effect of the edge of the state space. Alternatively, one could assume  $s(\cdot|\omega_i^R)$  is binomial, i.e., each unit of R&D stock is successfully incorporated with the same probability ( $q_n = \rho$  for all  $n$ ). However, if one uses the binomial specification, then one must assume a much lower success probability in order to successfully mitigate the effects of the edge of the state space. This lower success probability weakens investment incentives on the interior of the state space much more than the *generalized binomial* specification described above. In this respect, the *generalized binomial* specification is similar in spirit to the approach taken in Pakes and McGuire (1994), Gowrisankaran and Town (1997), and Borkovsky et al. (2012); see footnote 15.

#### A.4 A benchmark model with no version releases

In this section, I present the benchmark model that is discussed in Section 5. The benchmark model is a simplified version of the dynamic model, in which firms cannot stockpile R&D and accordingly do not make version release decisions. Specifically, I assume that when a firm achieves an R&D success, it is immediately incorporated into its product with certainty and at no additional cost.

The benchmark model is nested in the model in Section 2; to derive the benchmark model, one simply sets the release cost to zero ( $G_l = G_u = 0$ ) and the probability that a firm successfully incorporates its entire R&D stock when releasing a new version to one—i.e.,  $s(\omega_i^R | \omega_i^R) = 1$ . In this section, by incorporating these assumptions, I present the benchmark model and derive the Bellman equation and optimality conditions.

In the model in Section 2, firms make version release decisions in subperiod 1 and R&D investment decisions in subperiod 2. For the purposes of the benchmark model, it is convenient to regard a period as beginning with subperiod 2 and concluding at the end of the subsequent subperiod 1.<sup>53</sup> It follows that firms make R&D investment decisions (in step 5) and then incorporate the successful R&D outcomes (in the subsequent step 4) within the same period. Hence, because the R&D outcomes need not be tracked from period to period, the industry state is simply the product market state  $\omega^m$ .

Because of this reframing, I define  $\omega^m$  as the industry state at the beginning of subperiod 2 and  $V_i(\omega^m)$  as the expected net present value of all future cash flows to firm  $i$  in state  $\omega^m$  at the beginning of subperiod 2. I let  $\omega^{m'}$  be the industry state at the end of the subsequent subperiod 1. It follows that the law of motion is

$$\omega^{m'} = \omega^m + v_1 - v_2,$$


where  $v_i = 1$  if firm  $i$ 's R&D investment is successful and  $v_i = 0$  otherwise.

As in Section 2.2, because I solve for a symmetric equilibrium, it suffices to solve for the investment optimality condition of only one firm; hence I hereafter restrict attention to firm 1. Firm 1's value function  $V_1 : \Omega_d^m \rightarrow \mathbb{R}$  is defined recursively by the solution to the Bellman equation

$$V_1(\omega^m) = \max_{x_1 \geq 0} -x_1 + \beta \left\{ \pi_1(\omega^m) + \frac{\alpha x_1}{1 + \alpha x_1} Z_1(1 | \omega^m) + \frac{1}{1 + \alpha x_1} Z_1(0 | \omega^m) \right\}, \quad (10)$$

<sup>53</sup>Reframing the timing of the game does not change the model. Because the model is infinitely repeated, one can define the value function and accordingly derive the Bellman equation from the perspective of any step in the timing of a period—as long as one applies the discount factor  $\beta$  at the beginning of the actual period. Therefore, when deriving the Bellman equation, I apply the discount factor  $\beta$  at the beginning of subperiod 1.





**Fig. 9** Equilibrium price  $p_1(\omega^m)$  (left panel), market share (middle panel), and profit  $\pi_1(\omega^m)$  (right panel) for  $\sigma = 0.1, 2,$  and  $4$

where

$$Z_1(v_1|\omega^m) = \frac{\alpha x_2(\omega^m)}{1 + \alpha x_2(\omega^m)} V_1(\max(\omega^m + v_1 - 1, -L^m)) + \frac{1}{1 + \alpha x_2(\omega^m)} V_1(\min(\omega^m + v_1, L^m))$$

is firm 1’s expected value in industry state  $\omega^m$  conditional on an R&D investment success ( $v_1 = 1$ ) or failure ( $v_1 = 0$ ); and  $x_2(\omega^m)$  is the R&D investment of firm 2 in industry state  $\omega^m$ . The min and max operators merely enforce the bounds of the state space.


Solving the maximization problem on the right-hand side of the Bellman Eq. 10, I obtain the following optimality condition for firm 1’s R&D investment  $x_1(\omega^m)$ :

$$x_1(\omega^m) = \max \left\{ 0, \frac{-1 + \sqrt{\beta\alpha [Z_1(1|\omega^m) - Z_1(0|\omega^m)]}}{\alpha} \right\}. \tag{11}$$

if  $Z_1(1|\omega^m) - Z_1(0|\omega^m) \geq 0$ , and  $x_1(\omega^m) = 0$  otherwise. (Because I solve for a symmetric equilibrium,  $x_2(\omega^m) = x_1(-\omega^m)$ ). Solving for an equilibrium for a particular parameterization of the model amounts to finding a value function  $V_1(\cdot)$  and policy functions  $x_1(\cdot)$  that satisfy the Bellman Eq. 10 and the optimality condition 11 for all industry states  $\omega^m \in \Omega_d^m$ .

### A.5 The equilibrium correspondence

I thoroughly explore the effects of horizontal product differentiation on equilibrium behavior by computing equilibria for the baseline parameterization with a wide range of  $\sigma$  values, in particular,  $\sigma \in \{0.1, 0.2, \dots, 8\}$ . In Fig. 10, I summarize each equilibrium computed using the size of the leader’s expected lead after 30 periods (the industry’s expected lifespan),  $L^{30}$ . Figure 10 shows that I have found several regions



**Fig. 10** Equilibrium correspondence

in which there are multiple equilibria, including one region (around  $\sigma = 2$ ) in which there are up to four equilibria.<sup>54</sup>

When the degree of horizontal differentiation is sufficiently low ( $\sigma \leq 0.7$ ), the equilibria are qualitatively similar to the preemptive equilibrium in Fig. 3. At intermediate levels of horizontal differentiation ( $0.8 \leq \sigma \leq 2.6$ ), the equilibria are qualitatively similar to the accommodative equilibrium in Fig. 5. Finally, when the degree of horizontal differentiation is sufficiently high ( $\sigma \geq 2.7$ ), I have found up to three qualitatively different equilibria resembling those presented in Fig. 7.

By analyzing the four equilibria computed for the  $\sigma = 2$  parameterization, I am able to provide a deeper understanding of the accommodative equilibrium. Figure 5 shows that in the  $\omega^m = 0$  cross-section of the accommodative equilibrium, accommodative trenches arise at R&D stock levels of 1 and 4. This raises the following question: why do accommodative trenches arise at some levels of R&D stock and not others? In this case, is there something particularly special about R&D stock levels of 1 and 4? (I will show that the answer is “no”.)

It is difficult to discern the four equilibria for the  $\sigma = 2$  parameterization in Fig. 10 because the corresponding values of  $L^{30}$  are very similar: 9.5293, 9.8002,

<sup>54</sup>While I have succeeded in computing at least one equilibrium for each parameterization, there is no guarantee that all equilibria have been found.

9.8421, and 9.8816. One of these equilibria ( $L^{30} = 9.8421$ ) is the accommodative equilibrium presented in Fig. 5. The other three equilibria are qualitatively similar, the only difference being that the accommodative trenches arise at different R&D stock levels. This shows that for a given parameterization that admits accommodative equilibria, there can be several such equilibria, differing only in the locations of the accommodative trenches.

### A.6 An alternative model of product market competition: Hotelling with heterogeneous vertical qualities

The Logit demand model presented in Section 2.1 can be reinterpreted as an address model (Anderson et al. 1992); i.e., each consumer’s individual preferences are reflected by her fixed  $(\epsilon_1, \epsilon_2)$  values (her “location”), and the joint distribution of  $(\epsilon_1, \epsilon_2)$  reflects the distribution of taste heterogeneity among consumers. In this section, I present an alternative model of static price competition that incorporates both horizontal and vertical differentiation and that is rooted in a different address model—specifically, a Hotelling (1929) “linear city” model in which firms have fixed locations (at opposite ends of the Hotelling line) and different vertical qualities. I show that this model’s equilibrium profit function is qualitatively similar to the profit function of the Logit period game presented in Section 2.1. I then show that the equilibria of the corresponding dynamic model (which includes the Hotelling model as its period game) are qualitatively similar to the equilibria of the dynamic model presented in the paper.

**The model** A unit mass of consumers is uniformly distributed over the unit interval  $[0, 1]$ . Firm 0 (1) is located at  $x = 0$  ( $x = 1$ ) and sells a product of quality  $\omega_0^m$  ( $\omega_1^m$ ). Let  $\omega^m \equiv \omega_1^m - \omega_0^m$  denote the size of firm 1’s quality advantage. Each firm faces a constant marginal cost of production of  $c > 0$ . Let  $p_0$  and  $p_1$  denote the firms’ respective prices. A consumer incurs transportation cost  $t$  per unit of distance. I assume that the market is covered, i.e., each consumer buys one unit of one product. A consumer located at  $x \in [0, 1]$  receives utility  $U_0(x) = \omega_0^m - p_0 - tx$  from consuming firm 0’s product and  $U_1(x) = \omega_1^m - p_1 - t(1 - x)$  from consuming firm 1’s product. Each firm sets price to maximize profits.

**The equilibrium profit function** It is straightforward to show that the consumer who is indifferent between the products of firms 1 and 2 is located at  $\bar{x} = \frac{t - \omega^m + p_1 - p_0}{2t}$ . It follows that the firms face demand functions

$$D_0(p_0, p_1) = \frac{t - \omega^m + (p_1 - p_0)}{2t},$$

and

$$D_1(p_0, p_1) = \frac{t + \omega^m - (p_1 - p_0)}{2t}.$$

It is straightforward to solve each firm’s profit maximization problem and, thereafter, to solve for the equilibrium price, demand, and profit functions. Firms are symmetric,

hence I will restrict attention to firm 1. Firms 1's equilibrium price function is

$$p_1^*(\omega^m, t, c) = c + t + \frac{\omega^m}{3}. \quad (12)$$

Firm 1's equilibrium demand function is

$$D_1^*(\omega^m, t) = \frac{1}{2} + \frac{\omega^m}{6t}. \quad (13)$$

Firm 1's equilibrium profit function is

$$\pi_1^*(\omega^m, t) = \frac{t}{2} + \frac{\omega^m}{3} + \frac{(\omega^m)^2}{18t}. \quad (14)$$


I take the standard approach of deriving conditions on parameter values that ensure that the equilibrium prices  $p_0^*$  and  $p_1^*$  constitute an interior solution, i.e.,  $D_0(p_0^*, p_1^*) \in [0, 1]$  and accordingly  $D_1(p_0^*, p_1^*) \in [0, 1]$ . I find that assuming  $t \geq \frac{|\omega^m|}{3}$  guarantees an interior solution.

The profit function 14 is qualitatively similar to the profit function of the Logit period game presented in the paper. First, profit is a function of only the difference between firms' respective product qualities,  $\omega^m$ —not their absolute product qualities  $\omega_0^m$  and  $\omega_1^m$ . This follows directly from the assumptions (made in both models) that there is no outside good and that utility is linear in quality.

Second, the profit function is convex in the size of a firm's lead. This shows that the convexity of the Logit model's profit function (see Fig. 1) is not simply an artifact of the exponential nature of the Logit model. Rather, as the Hotelling model demonstrates, the convexity reflects the fact that an increase in a firm's quality advantage gives rise to both a higher equilibrium price (see Eq. 12) and—despite the higher price—a higher quantity demanded (see Eq. 13).

Third, an increase in the degree of horizontal differentiation (as reflected by the transportation cost  $t$  in the Hotelling model) has a qualitatively similar impact on profit in both models. To illustrate, I assume  $\omega^m \in [-20, 20]$  and present the profit functions for  $t = \frac{20}{3}$  and  $t = 17$  in Fig. 11. Comparing these profit functions to the profit functions presented in Fig. 1, we see that in both models, an increase in the degree of horizontal differentiation tends to increase profits and makes the profit function less convex.

As explained above, because I make the standard assumption that the equilibrium prices constitute an interior solution, I impose the condition  $t \geq \frac{|\omega^m|}{3}$ . When  $\omega^m \in [-20, 20]$ , I set  $t \geq \frac{20}{3}$  so as to ensure that this condition is satisfied for all  $\omega^m$  values. Because of this condition, I can reduce the degree of horizontal differentiation in the Hotelling model only so much. This is reflected in the profit function for  $t = \frac{20}{3}$  presented in Fig. 11, which is qualitatively similar to the one presented in Fig. 1 for  $\sigma = 2$ . It is characterized by moderate horizontal differentiation, which is reflected by the fact that the leader's (follower's) profit increases (declines) relatively slowly as the leader's lead increases. Because of the  $t \geq \frac{20}{3}$  condition, the model does not admit an interior solution that gives rise to an equilibrium profit function similar to the one presented in Fig. 1 for  $\sigma = 0.1$ . Rather, a lower degree of horizontal differentiation can only be incorporated if one allows for corner solutions in which one

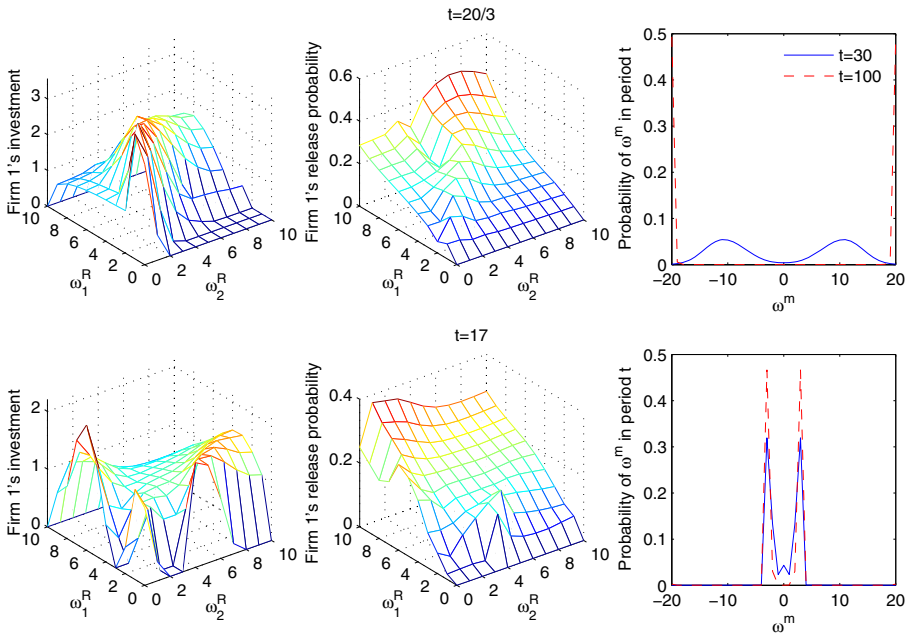


**Fig. 11** Hotelling model profit function

firm captures the entire market (which gives rise to a kinked profit function). This highlights the benefit of using the Logit model of product market competition presented in Section 2, i.e., it is straightforward to generate wide variation in the degree of horizontal differentiation because one can easily solve for an interior solution for any  $\sigma > 0$ .

**Equilibria of the dynamic model** As explained in Section 2, the prices firms set in the (static) product market game do not impact state-to-state transitions in the dynamic model. Hence the period game profit function can be treated as a “primitive” of the dynamic model; i.e., for the purposes of the dynamic model, it is as if the period game profit function is exogenous (see p. 1892 of Doraszelski and Pakes 2007). It follows that two *different* models of product market competition that give rise to qualitatively similar profit functions should yield qualitatively similar equilibria of the dynamic model. To verify this for the Hotelling model, I compute equilibria of the dynamic model for the two profit functions presented in Fig. 11 (holding all dynamic model parameters equal to their baseline values) and present them in Fig. 12.

Although they differ in terms of their scales, the profit functions for  $t = \frac{20}{3}$  and  $t = 17$  in Fig. 11 are qualitatively similar to those for  $\sigma = 2$  and  $\sigma = 4$  in Fig. 1. Figure 12 shows that the corresponding equilibria of the dynamic model are similar



**Fig. 12** Equilibria of dynamic model with Hotelling period game for  $t = \frac{20}{3}$  (top row) and  $t = 17$  (bottom row).  $\omega^m = 0$  cross-sections of policy functions for R&D investment  $x_1(\omega)$  (left panels) and release probability  $r_1(\omega)$  (middle panels). Transient distributions  $\mu_t^m(\omega^m)$  for  $t = 30$  and  $t = 100$  (right panels)

as well. The dynamic model equilibrium for  $t = \frac{20}{3}$  in the top row of Fig. 12 is qualitatively similar to the one for  $\sigma = 2$  in Fig. 5 in that they are both characterized by moderate preemption and phases of accommodation. The dynamic model equilibrium for  $t = 17$  in the bottom row of Fig. 12 is qualitatively similar to the one for  $\sigma = 4$  in second row of Fig. 7 in that they are both characterized by asymmetric R&D wars. (I present only one equilibrium for the  $t = 17$  parameterization because I have not searched for multiplicity.) As explained above, the Hotelling model does not admit an interior solution that gives rise to an equilibrium profit function similar to the one presented in Fig. 1 for  $\sigma = 0.1$ ; hence, the corresponding dynamic model does not admit an equilibrium characterized only by intense preemption races, like the one presented in Fig. 3.

### A.7 Alternative versions of the dynamic model

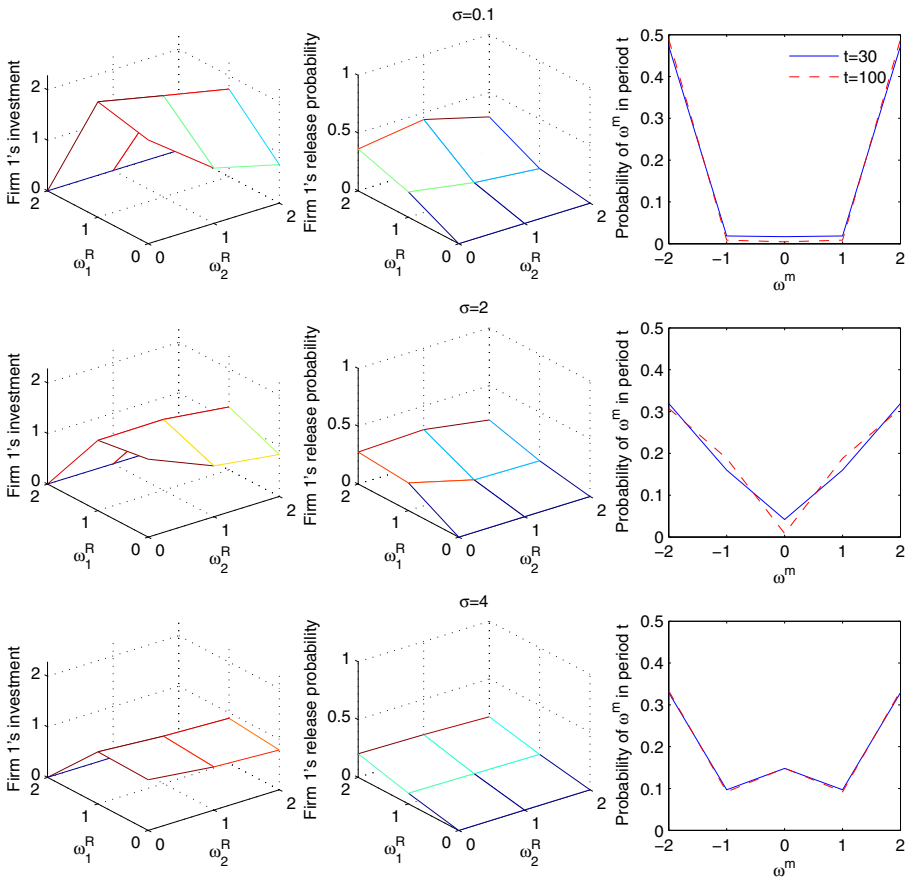
In this section, I present results for two alternative versions of the dynamic model. The results and the related discussions are useful in two respects. First, they help readers understand why some of the key assumptions made in the dynamic model are needed. Second, they help readers better understand the intuitions that underly the results presented in the paper.

First, I present results for a version of the model with a much smaller state space. The results shows that a model with a small state space is not rich enough to admit

the behaviors described in Section 6. Second, I present results for a version of the model in which there is no uncertainty in the version release process; i.e., when a firm releases a new version, it always succeeds in incorporating its entire R&D stock. I use these results to explain the impact of uncertainty in the version release process on firm behavior. I then show that the equilibria of this version of the model are qualitatively similar to the corresponding equilibria presented in the paper.

*A.7.1 A smaller state space*

In this section, I explore whether the equilibrium behaviors described in Section 6 would arise in a model with a much smaller state space. I make the state space as small as possible while still allowing for the accumulation of R&D successes; i.e.,



**Fig. 13** Model with a smaller state space. Equilibria for  $\sigma = 0.1$  (top row),  $\sigma = 2$  (middle row), and  $\sigma = 4$  (bottom row).  $\omega^m = 0$  cross-sections of policy functions for R&D investment  $x_1(\omega)$  (left panels) and release probability  $r_1(\omega)$  (middle panels). Transient distributions  $\mu_t^m(\omega^m)$  for  $t = 30$  and  $t = 100$  (right panels)

I set  $L^R = 2$ , which implies that  $\omega_i^R \in \{0, 1, 2\}$ , and  $L^m = 2$ , which implies that  $\omega^m \in \{-2, \dots, 2\}$ . So, each firm can possess either 0, 1, or 2 units of R&D stock and can hold a product market lead of either 0, 1 or 2 units.

Figure 13 presents equilibria for the baseline parameterization with  $\sigma = 0.1, 2$ , and 4, which correspond to the equilibria presented in Figs. 3, 5 and 7. (The corresponding period profit functions are simply the profit functions in Fig. 1 confined to  $\omega^m \in \{-2, \dots, 2\}$ . I present only one equilibrium for the  $\sigma = 4$  parameterization because I have not searched for multiplicity.) Fig. 13 shows that the model with a small state space does not admit any of the behaviors described in Section 6: preemption races, phases of accommodation, and asymmetric R&D wars.

The discussion of these behaviors in Section 6 suggests that they arise only if the state space is much larger because the model then provides a sufficiently rich reflection of how changes in the competitive positioning of firms impact their period profits and accordingly their investment and updating incentives. In other words, the dynamic incentives that firms face can differ drastically depending on whether the leader's advantage is non-existent, small, moderate, large, very large etc.

Figure 13 presents equilibria for a "smaller" model. Alternatively, one could consider a "coarser" model, i.e., one with a coarser discretization of the state space. Instead of restricting attention to the product market states  $\omega^m \in \{-2, -1, 0, 1, 2\}$ , one could restrict attention to the states  $\omega^m \in \{-20, -10, 0, 10, 20\}$ . I have found that equilibria of a coarser model are qualitatively similar to those presented in Fig. 13 (just, expectedly, with higher investment levels and updating probabilities) and accordingly do not admit the behaviors described in Section 6. These results are available upon request.

### A.7.2 No uncertainty in version releases

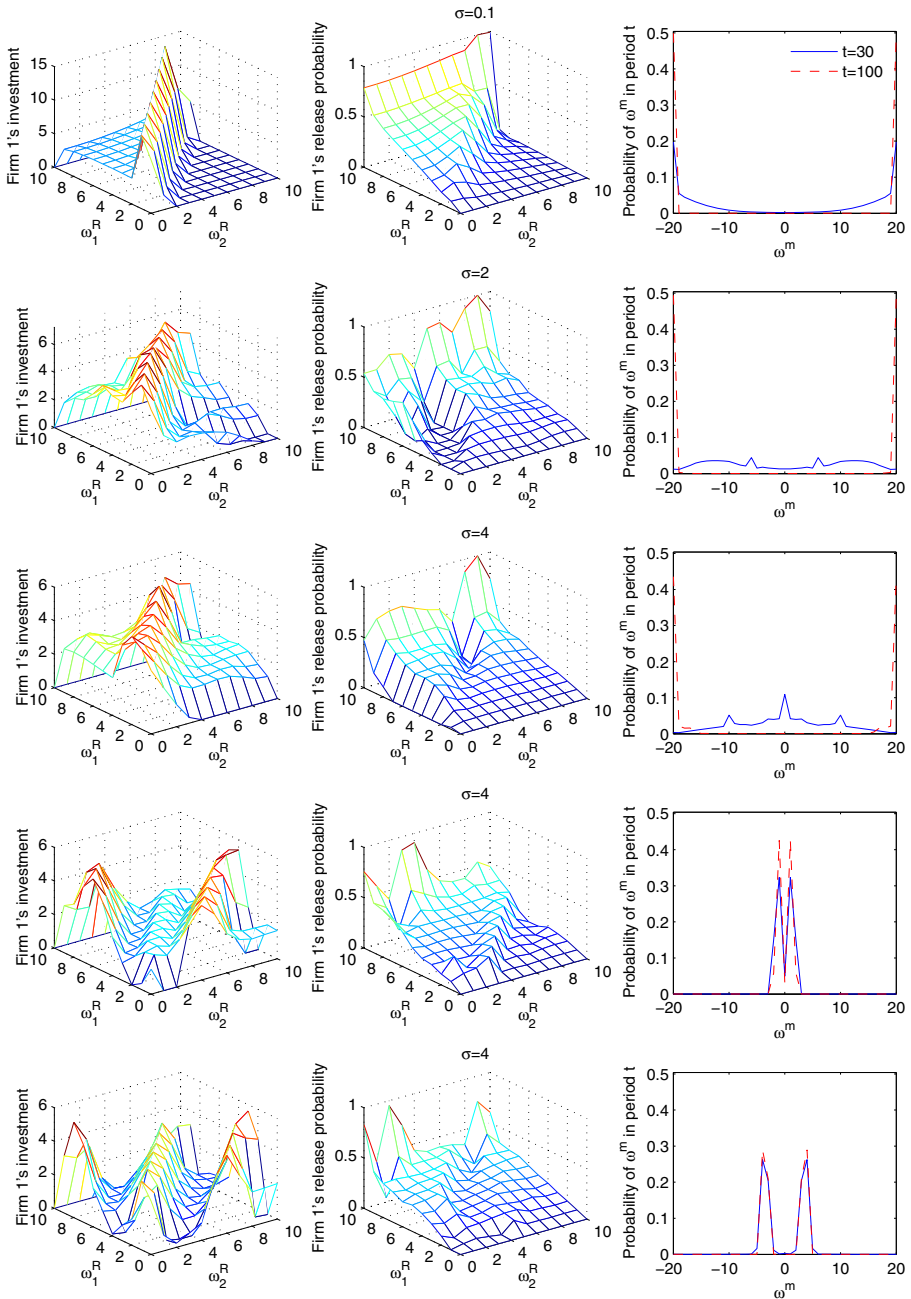
In this section, I explore a version of the model in which there is no uncertainty in the version release process. That is, whenever a firm releases a new version, it successfully incorporates its entire R&D stock.<sup>55</sup> Figure 14 presents equilibria for the baseline parameterization with  $\sigma = 0.1, 2$ , and 4, which correspond to the equilibria presented in Figs. 3, 5 and 7.

I first explain the effect of uncertainty in the version release process on firm behavior. I then explain that aside from this effect, these equilibria are qualitatively similar to the corresponding equilibria presented in the paper. Finally, in Section A.8, I use these results to explain how this uncertainty mitigates the effect of the edge of the state space.

The top row of Fig. 14 presents the equilibrium for a low degree of horizontal differentiation ( $\sigma = 0.1$ ). It is similar to the corresponding equilibrium in Fig. 3 in that it is characterized by preemption in terms of both R&D investment and release probabilities. However, there are two qualitative differences. First, the R&D preemption is more intense in the sense that firms invest more heavily in R&D during a preemption

<sup>55</sup>This simply entails setting  $s(\omega_i^R|\omega_i^R) = 1$  for all  $\omega_i^R \in \Omega^R$ .





**Fig. 14** Model with no uncertainty in version releases. Five equilibria, for  $\sigma = 0.1$  (top row),  $\sigma = 2$  (second row), and  $\sigma = 4$  (rows 3-5).  $\omega^m = 0$  cross-sections of policy functions for R&D investment  $x_1(\omega)$  (left panels) and release probability  $r_1(\omega)$  (middle panels). Transient distributions  $\mu_t^m(\omega^m)$  for  $t = 30$  and  $t = 100$  (right panels)

race and the preemption race comes to an end more quickly. Second, the preemption in terms of release probabilities is milder.

I will first explain why eliminating uncertainty from the version release process makes R&D preemption more intense. If a firm's entire R&D stock is successfully incorporated when it releases a new version, then R&D investment becomes a more effective tool for enhancing product quality and, accordingly, for building a lead over a rival in hopes of inducing that rival to give up. This gives rise to stronger investment incentives when neither firm has too large a lead. The preemption race ends more quickly because a smaller R&D stock advantage is required to induce a rival to give up, simply because the rival recognizes that that R&D stock advantage will necessarily be translated into a quality advantage in the product market.

To understand why eliminating uncertainty from the version release process gives rise to milder preemption in terms of release probabilities, recall that firms engage in preemption for purely strategic reasons and that the strategic benefit of a version release stems directly from the fact that it is uncertain (see Section 6.1.1). In the absence of uncertainty, while a version release does allow a firm to gain a quality advantage in the product market, it compromises its ability to gain a strategic advantage.<sup>56</sup>

Except for the effects of eliminating uncertainty that are described above, the equilibria presented in rows 2-5 of Fig. 14 are qualitatively similar to the corresponding equilibria in Section 6. The equilibria in the second and third rows of Fig. 14 are characterized by preemption and phases of accommodation just like the ones in Fig. 5 and the top row of Fig. 7. The equilibrium in the fourth row of Fig. 14 is characterized by asymmetric R&D wars (as well as some mild R&D preemption stemming from the first effect described above) just like the equilibrium in the second row of Fig. 7. The equilibrium in the bottom row of Fig. 14 is characterized by mild R&D preemption and asymmetric R&D wars just like the equilibrium in the bottom row of Fig. 7. Accordingly, the short- and long-run industry structures presented in the right column of Fig. 14 are qualitatively similar to the industry structures of the corresponding equilibria presented in Section 6.

### A.8 The effect of the edge of the state space

In this section, I discuss an example that shows how introducing uncertainty into the version release process mitigates the effects of the edge of the state space. Consider the equilibrium in the third row of Fig. 14. In industry state  $(0, 10, 9)$ , firm 1 holds a small R&D stock lead over firm 2, but cannot increase this lead further because it has already achieved the highest possible R&D stock level. Because its only recourse is to release a new version, it does so with very high probability (specifically, probability 1). This gives rise to perverse R&D investment incentives for firm 2; in industry states

<sup>56</sup>This raises the question of why the release probabilities in the top row of Fig. 14 are characterized by any preemption whatsoever. The reason is that there is also uncertainty over the cost of releasing a new version. (Recall that this uncertainty is required to guarantee the existence of an equilibrium in pure strategies.) Hence, a firm that releases a new version can gain a strategic advantage over a rival because to catch up the rival would need to draw a release cost that is low enough to make updating optimal.

near (0, 10, 9), firm 2 invests very heavily in R&D in hopes of avoiding industry state (0, 10, 9).<sup>57</sup> Firm 1 best responds by investing heavily and updating with high probability in nearby states.

Comparing this equilibrium with the corresponding equilibrium in the top row of Fig. 7 shows that incorporating uncertainty into the version release process mitigates the perverse investment and updating incentives caused by the edge of the state space. Specifically, in the presence of uncertainty, the incentives that firm 1 faces in industry state (0, 10, 9) are not as different from those it faces in neighboring state (0, 9, 9) because in both states it understands that it is unlikely to successfully incorporate its entire R&D stock when it releases a new version.

## References

- Aguirregabiria, V., & Nevo, A. (2013). Recent developments in empirical IO: dynamic demand and dynamic games. In D. Acemoglu, M. Arellano, & E. Dekel (Eds.) *Advances in economics and econometrics, 10th world congress* (Vol. 3, pp. 53–122). Cambridge: Cambridge University Press.
- Aizcorbe, A. (2005). Moore's Law, competition, and Intel's productivity in the mid-1990s. *American Economic Review Papers and Proceedings*, 95(2), 305–308.
- Alfonsi, S., Fahy, U., & Brown, K. (2010). *Battle for baby's bottom: diaper wars heat up*. ABC News.
- Anderson, S., de Palma, A., & Thisse, J. (1992). *Discrete choice theory of product differentiation*. Cambridge: MIT Press.
- Aoki, R. (1991). R&D competition for product innovation: an endless race. *American Economic Review*, 81(2), 252–256.
- Bajari, P., Hong, H., & Nekipelov, D. (2013). Game theory and econometrics: a survey of some recent research. In D. Acemoglu, M. Arellano, & E. Dekel (Eds.) *Advances in economics and econometrics, 10th world congress* (Vol. 3, pp. 3–52). Cambridge: Cambridge University Press.
- Baker, M., Graham, P., Harker, D., & Harker, M. (2000). *Marketing: managerial foundations*. South Yarra: Macmillan Publishers.
- Besanko, D., Doraszelski, U., Kryukov, Y., & Satterthwaite, M. (2010). Learning-by-doing, organizational forgetting, and industry dynamics. *Econometrica*, 78(2), 453–508.
- Bhaskaran, S., & Ramachandran, K. (2011). Managing technology selection and development risk in competitive environments. *Production & Operations Management*, 20(4), 541–555.
- Borkovsky, R., Doraszelski, U., & Kryukov, S. (2012). A dynamic quality ladder duopoly with entry and exit: exploring the equilibrium correspondence using the homotopy method. *Quantitative Marketing and Economics*, 10(2), 197–229.
- Borkovsky, R., Goldfarb, A., Haviv, A., & Moorthy, S. (2017). Measuring and understanding brand value in a dynamic model of brand management. *Marketing Science*, 36(4), 471–499.
- Bradsher, K. (2010). *China drawing high-tech research from USA*. The New York Times.
- Bresnahan, T. (2001). *Network effects and microsoft, stanford institute for economic policy research discussion paper No. 00-51*. Stanford: Stanford University.
- Budd, C., Harris, C., & Vickers, J. (1993). A model of the evolution of duopoly: does the asymmetry between firms tend to increase or decrease? *Review of Economic Studies*, 60(3), 543–573.
- Caplin, A., & Nalebuff, B. (1991). Aggregation and imperfect competition: on the existence of equilibrium. *Econometrica*, 59(1), 26–59.
- Chain Drug Review (2010). It's dueling shavers as Gillette, Schick launches catch on.
- Dahremoller, C., & Fels, M. (2015). Product lines, product design, and limited attention. *Journal of Economic Behavior and Organization*, 119, 437–456.

<sup>57</sup>Recall that the  $\omega = 0$  cross-sections of firm 2's policy functions are the transposes of those presented for firm 1.

- Doraszelski, U., & Markovich, S. (2007). Advertising dynamics and competitive advantage. *RAND Journal of Economics*, 38(3), 1–36.
- Doraszelski, U., & Pakes, A. (2007). A framework for applied dynamic analysis in IO. In M. Armstrong, & R. Porter (Eds.) *Handbook of industrial organization* (Vol. 3 pp. 1887–1966). North-Holland, Amsterdam.
- Doraszelski, U., & Satterthwaite, M. (2010). Computable Markov-perfect industry dynamics. *RAND Journal of Economics*, 41(2), 215–243.
- Dubé, J., Hitsch, G., & Chintagunta, P. (2010). Tipping and concentration in markets with indirect network effects. *Marketing Science*, 29(2), 216–249.
- Dyer, D., Dalzell, F., & Olegario, R. (2004). *Rising tide: lessons from 165 years of brand building at procter & gamble*. Boston: Harvard Business School Press.
- Elzinga, K., & Mills, D. (1996). Innovation and entry in the US disposable diaper industry. *Industrial and Corporate Change*, 5(3), 791–812.
- Ericson, R., & Pakes, A. (1995). Markov-perfect industry dynamics: a framework for empirical work. *Review of Economic Studies*, 62, 53–82.
- Fried, I. (2009). *Microsoft launches IE8 with a smile*. CNET News.
- Fudenberg, D., Gilbert, R., Stiglitz, J., & Tirole, J. (1983). Preemption, leapfrogging and competition in patent races. *European Economic Review*, 22, 3–31.
- Goettler, R., & Gordon, B. (2011). Does AMD spur Intel to innovate more? *Journal of Political Economy*, 119(6), 1141–1200.
- Goettler, R., & Gordon, B. (2014). Competition and product innovation in dynamic oligopoly. *Quantitative Marketing and Economics*, 12(1), 1–42.
- Gordon, B. (2009). A dynamic model of consumer replacement cycles in the PC processor industry. *Marketing Science*, 28(5), 846–867.
- Gowrisankaran, G., & Town, R. (1997). Dynamic equilibrium in the hospital industry. *RAND Journal of Economics*, 6(1), 45–74.
- Grossman, G., & Shapiro, C. (1987). Dynamic R&D competition. *Economic Journal*, 97(386), 372–387.
- Hansen, E. (2003). *Microsoft to abandon standalone IE*. CNET News.
- Harris, C. (1991). *Dynamic competition for market share: an undiscounted model*. Working paper, Nuffield College, Oxford University: Oxford.
- Harris, C., & Vickers, J. (1985). Perfect equilibrium in a model of a race. *Review of Economic Studies*, 52(2), 193–209.
- Harris, C., & Vickers, J. (1987). Racing with uncertainty. *Review of Economic Studies*, 54(1), 1–21.
- Harsanyi, J. (1973). Games with randomly disturbed payoffs: a new rationale for mixed-strategy equilibrium points. *International Journal of Game Theory*, 2(1), 1–23.
- Hartmann, W., & Nair, H. (2010). Retail competition and the dynamics of demand for tied goods. *Marketing Science*, 29(2), 366–386.
- Hitsch, G. (2006). An empirical model of optimal dynamic product launch and exit under demand uncertainty. *Marketing Science*, 25(1), 25–50.
- Hoover, N. (2006). The fight is on: Internet Explorer 7.0 takes on Firefox 2.0 in a face-off that's reminiscent of the browser battle of the '90s. This time, it's Microsoft's fight to lose. *Information Week*.
- Hörner, J. (2004). A perpetual race to stay ahead. *Review of Economic Studies*, 71(4), 1065–1088.
- Hotelling, H. (1929). Stability in competition. *Economic Journal*, 39, 41–58.
- Hout, T., & Ghemawat, P. (2010). China vs the world: whose technology is it? *Harvard Business Review*, 88(12), 94–103.
- Jansen, J. (2010). Strategic information disclosure and competition for an imperfectly protected innovation. *Journal of Industrial Economics*, 58(2), 349–372.
- Judd, K. (1998). *Numerical methods in economics*. Cambridge: MIT Press.
- Kamien, M., & Schwartz, N. (1972). Timing of innovations under rivalry. *Econometrica*, 40(1), 43–60.
- Lee, R. (2013). Vertical integration and exclusivity in platform and two-sided markets. *American Economic Review*, 103(7), 2960–3000.
- Levitt, T. (1960). Marketing myopia. *Harvard Business Review*, 38(4), 45–56.
- Lippman, S., & McCardle, K. (1987). Dropout behavior in R&D races with learning. *RAND Journal of Economics*, 18(2), 287–295.
- Lippman, S., & McCardle, K. (1988). Preemption in R&D races. *European Economic Review*, 32, 1661–1669.

- Markoff, J. (2001). *USA vs. Microsoft: news analysis; making a judgement on a moving target*. The New York Times.
- Markovich, S. (2008). Snowball: a dynamic oligopoly model with network externalities. *Journal of Economic Dynamics and Control*, 32(3), 909–938.
- Markovich, S., & Moenius, J. (2009). Winning while losing: competition dynamics in the presence of indirect network effects. *International Journal of Industrial Organization*, 27(3), 346–357.
- McCoy, M. (2001). What's that stuff? Fluoride. *Science & Technology*, 79(16), 42.
- Moorthy, S. (1988). Product and price competition in a duopoly. *Marketing Science*, 7(2), 141–168.
- Moorthy, S., & Png, P. (1992). Market segmentation, cannibalization, and the timing of product introductions. *Management Science*, 38(3), 345–359.
- Morgan, L., Morgan, R., & Moore, W. (2001). Quality and time-to-market trade-offs when there are multiple product generations. *Manufacturing & Service Operations Management*, 3(2), 89–104.
- Nair, H., Chintagunta, P., & Dubé, J. (2004). Empirical analysis of indirect network effects in the market for personal digital assistants. *Quantitative Marketing and Economics*, 2(1), 23–58.
- Neven, D., & Thisse, J. (1990). On quality and variety competition. In J. Gabszewicz, J. Richard, & L. Wolsey (Eds.) *Economic decision-making: games, econometrics and optimisation* (pp. 175–199). Amsterdam: Elsevier.
- Ofek, E., & Sarvary, M. (2003). R&D, marketing and the success of next-generation products. *Marketing Science*, 2(3), 355–370.
- Pakes, A., & McGuire, P. (1994). Computing Markov-perfect Nash equilibria: numerical implications of a dynamic differentiated product model. *RAND Journal of Economics*, 25(4), 555–589.
- Parry, M. (2001). *Crest toothpaste: the innovation challenge*. Darden Business Publishing.
- Parry, M., & Jones, M. (2001). *Pampers: the disposable diaper war (A)*. Darden Business Publishing.
- Prendergrast, M. (2000). *For God, Country & Coca-Cola: the definitive history of the great American soft drink and the company that makes it*. New York: Basic Books.
- Qi, S. (2013). The impact of advertising regulation on industry: the cigarette advertising ban of 1971. *RAND Journal of Economics*, 44(2), 215–248.
- Ramachandran, K., & Krishnan, V. (2008). Design architecture and introduction timing for rapidly improving industrial products. *Manufacturing & Service Operations Management*, 10(1), 149–171.
- Richardson, A. (2010). *Innovation X: why a company's toughest problems are its greatest advantage*. San Francisco: Jossey Bass.
- Rietschel, R., & Fowler, J. (2008). *Fisher's contact dermatitis 6*. Hamilton: BC Decker Inc.
- Rivkin, S., & Sutherland, F. (2004). *The making of a name: the inside story of the brands we buy*. New York: Oxford University Press.
- Scherer, F. (1967). Research and development allocation under rivalry. *Quarterly Journal of Economics*, 81, 351–394.
- Shaked, A., & Sutton, J. (1982). Relaxing price competition through product differentiation. *Review of Economic Studies*, 49(1), 3–13.
- Sutton, J. (1991). *Sunk costs and market structure*. Cambridge: MIT Press.
- Tirole, J. (1988). *The theory of industrial organization*. Cambridge: MIT Press.
- West, C. (2001). *Competitive intelligence*. Basingstoke: Palgrave Macmillan.
- Wildstrom, S. (2003). *On beyond microsoft explorer*. Business Week.
- Wilson, L., & Norton, J. (1989). Optimal entry timing for a product line extension. *Marketing Science*, 8(1), 1–17.