Available online at www.sciencedirect.com

ScienceDirect

Review of Economic Dynamics 11 (2008) 386–412

Review of Economic Dynamics

www.elsevier.com/locate/red

Strategic experimentation and disruptive technological change ☆

Fabiano Schivardi [a,b,c,g], Martin Schneider [c,d,e,f,*]

[a] *University of Cagliari, Italy*
[b] *CRENoS, Italy*
[c] *CEPR, UK*
[d] *New York University, New York, USA*
[e] *Federal Reserve Bank of Minneapolis, MN, USA*
[f] *National Bureau of Economic Research, USA*
[g] *EIEF, Italy*

Received 19 May 2006; revised 21 September 2007

Available online 16 October 2007

Abstract

This paper studies the diffusion of a new technology that is brought to market while its potential is still uncertain. We consider a dynamic game in which an incumbent and a startup firm improve both a new and a rival old technology while learning about the relative potential of both technologies. The main findings are that (i) risk considerations make incumbents with higher market shares *more* likely to adopt the new technology and (ii) changes in market power are often preceded by a subpar performance of the new technology. We also show that introducing a better new technology or confronting a worse old technology may *hurt* the startup firm as its new technology is then adopted earlier by incumbents.

JEL classification: C73; C63; L13; D83; O31

Keywords: Oligopoly; Innovation; Learning; Dynamic games

1. Introduction

Many industries undergo long periods of incremental technological change, punctuated occasionally by *episodes of disruptive change*.[1] An episode of disruptive change features (i) startup firms who introduce a new design for a process or product, (ii) uncertainty, at least initially, about whether the new design can replace the currently dominant "old" design, and (iii) a period during which old and new designs compete. During this period, incumbents and startups race to incrementally improve their product performance and gain market share, while learning about the potential of the new design. A challenge for the incumbent is to decide whether and when to adopt the new design.

☆ The views expressed here are our own and do not necessarily reflect those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.
* Corresponding author at: Research Department, Federal Reserve Bank of Minneapolis, 90 Hennepin Ave., Minneapolis, MN 55480, USA.
E-mail addresses: fschivardi@unica.it (F. Schivardi), ms1927@nyu.edu (M. Schneider).
[1] The term was coined by Tushman and Anderson (1986, 1990), who provide an overview of many industry case studies.

1094-2025/$ – see front matter
doi:10.1016/j.red.2007.09.002

This paper models an episode of disruptive change as a dynamic investment game with learning and technology adoption. We emphasize two key properties of such episodes. First, in a strategic context, adopting an unproven new design *reduces* the risk faced by incumbents: competing with equal designs bounds the advantage that any one firm can achieve, whatever the absolute potential of the new design. This implies that incumbents who have a large market share that they want to protect (rather than expand) should be *more* willing to reduce risk by adopting. In contrast, as incumbents' market share slips, they have less to lose, become more willing to take risk and thus more reluctant to adopt. Second, learning assigns a special role to the early history of performance under the new design. In particular, a startup with a superior new design is more likely to gain market dominance if its design initially does not perform too well. This prediction is consistent with case study evidence on the cross section of episodes of disruptive change.

In our model, firm profits depend on a single state variable, relative product performance. Two firms—incumbent and startup—can invest to improve their performance by undertaking fixed size investment projects. In this respect, the model looks like a multi-stage patent race. The new feature is a choice of design: the *potential* of a design is identified with the probability that an investment project succeeds in improving performance; the design thus determines how quickly performance can be improved on average. The incumbent initially employs an old design of known potential. He may choose to adopt the new design of unknown potential that is employed by the startup. Firms learn about the potential of the new design in a Bayesian fashion by observing each other's investment and performance.

The incumbent's adoption decision is thus a choice between two races, a safer one with equal designs (that can be entered through adoption) and a riskier one with unequal designs (that is run before adoption). The incumbent must also take into account the option value of waiting, since the adoption decision is irreversible. Incumbents' choice of race depends on their market share, which shapes their attitude towards risk. In particular, market leaders face diminishing marginal gains from better performance and thus act in a risk averse manner.[2] They also perceive a smaller option value of waiting. In contrast, laggards with small market share face increasing marginal gains; they are risk-loving and perceive a larger option value.

Incumbents' optimal management of risk implies that the adoption decision is not based on expected potential alone. On the one hand, incumbents with a high market share may want to adopt the new design even if its expected potential is worse than that of the old design. This is because they prefer to reduce risk by moving to a race with equal designs. On the other hand, incumbents faced with a decline in market share become less and less likely to adopt: they prefer to "gamble for resurrection" by racing with unequal designs.

The model generates episodes of disruptive change of finite (but random) length. At the end of every episode, one of the two firms dominates the market. The market thus selects a new "dominant design," employed by the market leader. The final outcome, as well as the speed of diffusion, are determined by the *early history* of incremental innovations that occurs right after a new design has first been introduced. In particular, incumbents delay adoption if the evidence on the new design is initially unfavorable. If the startup subsequently catches up, the incumbent not only updates his belief, but also becomes more risk-loving as his market share declines. He thus becomes more reluctant to adopt as he prefers the risky race with unequal designs.

The model thus predicts that episodes in which there is late adoption and a change in market structure begin with slow initial performance growth under the new design. This result is consistent with evidence on the cross-section of episodes discussed by Tushman and Anderson (1990) and Christensen (1997). These authors document that "incumbent inertia"—a dramatic loss in incumbent market share accompanied by a failure to quickly adopt a new technology—is often preceded by a phase during which the new technology initially languishes in a market niche for some time.[3]

[2] Formally, firms in our model have risk neutral preferences over payoffs, as is standard in oligopoly models. The term "risk attitude" refers to the curvature of the profit function, which is convex in relative performance when the latter is low and concave when it is high. This interpretation is helpful when thinking about a choice between races, which changes the stochastic process for performance, but leaves the map from performance to payoffs unchanged.

[3] The hard disk drive industry, described in detail by Christensen (1993), provides a well-known example of this phenomenon. Over the period from the mid seventies to the mid eighties, the diameter of hard drives was reduced four times, from 14" to 11", 8", 5.25" and eventually to 3.5". Every time a smaller drive (that is, a new design) was introduced, always by an entrant firm, the industry witnessed at least 2 years of head-to-head competition between the new, smaller, drive and the dominant, larger, drive. Initially, it was not clear whether the lower storage capacity of smaller drives could ever be increased sufficiently to replace the larger drive. This was only determined over time as both designs were improved through incremental innovation. In the meantime the incumbent firm preferred to keep developing the larger drive, often ending up losing market leadership. Similar patterns have been documented for the excavator industry after World War II, when new hydraulically-actuated equipment was introduced

In the model, the startup's chance of winning is larger if the potential of the new design is not too high. Indeed, a higher potential of the new design has two effects. On the one hand, it generates a performance "drift" in the startup's favor *before* adoption, which allows the startup to catch up faster before adoption. On the other hand, it also tends to cut short the pre-adoption phase, since the new design generates more positive signals on average. Moreover, catching up *after* adoption is more difficult if the potential of the new design is very high. This is because a large probability of success of firms' investment projects means that investing firms typically move in lockstep. Indeed, the probability of catching up would be zero if both firms were to employ a design for which performance growth is deterministic. As a result, an increase in the potential of the new design can be good for the startup if the potential is low, but it becomes bad for the startup once the potential is high. The same forces imply that, if the potential of the new design is high, then it is better for the startup if the old design is actually *better*. This leads to a longer pre-adoption phase during which the startup can catch up, whereas it is very difficult to catch up after adoption if the new design is very good.

Our approach also emphasizes the possibility that competition selects the worse of the two designs. Two types of mistake can occur. On the one hand, a superior new design may not become dominant, because its early history is unfavorable and the incumbent retains market leadership without adopting. This type I mistake is reminiscent of simple armed bandit models: there is too little experimentation before the unproven new design is abandoned. On the other hand, our model also gives rise to type II mistakes: an inferior new design may be *preemptively adopted* by the incumbent in order to reduce the risk of racing with unequal designs. Since learning generally continues after adoption we may then end up in a situation where everybody agrees that the new design turned out to be inferior. As a general rule, both types of mistake are more likely when the incumbent initially has a larger performance advantage. This not only increases the incumbent's probability of driving out the startup, but also increases his propensity to adopt.

Most of our results are numerical—we compute Stationary Markov Perfect Equilibria (SMPE) of our dynamic game. Our computational algorithm has proven to be a reliable tool and may be of independent interest. SMPE present a challenge because they often cannot be a solved by simple value function iteration. In addition, we have to deal with both beliefs and performance as state variables.[4] We solve the first problem by using a combination of value function iteration and a "local guess-and-verify" procedure that deals with nonconvergence in problem areas of the state space. We address the dimensionality problem by exploiting the nonstationarity of the learning process, first approximating the value function in states where beliefs have almost converged, and then backtracking from there, thus avoiding costly iteration.

The paper is structured as follows. The remainder of this introduction reviews related literature. Section 2 presents the model. Section 3 characterizes firms' equilibrium strategies: here we show that an equilibrium produces an episode that ends in finite time and how the adoption decision depends on the state variables, beliefs and performance. Section 4 turns to simulation analysis of equilibria. Here we establish predictions that link the early history to the evolution of market structure. Section 5 concludes.

1.1. Related literature

There is a large body of work on dynamic investment races. Our paper differs from earlier work because it allows for learning and a choice of races for the incumbent (through the adoption decision). If these features were absent, our setup would reduce to a multistage investment game with a single state variable driving profits, as studied by Budd et al. (1993). In fact, if the potential of a new technology were known, then the subgame that firms in our model enter *after* adoption would be a discrete-time version of the game studied by these authors. Cabral (2003), Anderson and Cabral (in press) and Judd (2003) have considered multistage investment games in which firms choose the riskiness of investment projects at every stage. They show that firms who are behind in the race choose more risky strategies. This result is related to our finding that incumbents who have fallen behind want to stick with a more risky race, which means in our setting that they do not want to adopt the uncertain new design.

Delayed adoption has also been explained by various types of firm-specific adoption costs. In vintage capital models, it is assumed that (perfectly competitive) firms accumulate expertise—referred to as "human capital" by Chari and Hopenhayn (1991) or "experience" by Parente (1994)—that is tied to a particular technology. Adopting a

as an alternative to then dominant cable-actuated machines (Christensen, 1997), as well as for the cement, glass and mini-computer industries (Tushman and Anderson, 1990).

[4] Existing algorithms for investment games, such as Pakes and McGuire (1994), do not allow for learning and technology choice.

new technology entails the cost of losing some expertise. Similarly, in Jovanovic and Nyarko (1996), a single firm learns about how to best operate a technology. Adoption of a new technology creates uncertainty about how to operate that technology and firms need to regain expertise through learning-by-doing. One difference between these models and ours are that they assume the potential of every new technology to be known.[5] The early history of incremental innovations under a new technology therefore does not matter for diffusion. In addition, these models are not designed to make predictions about market structure.

A number of papers have studied adoption under uncertainty by a single firm. Jensen (1982) first modeled uncertainty as a reason for delayed adoption. In his model, a firm receives a sequence of costless signals about the potential of a new technology and optimally adopts when it becomes optimistic enough. In particular, a worse new technology might be adopted "by mistake" if the early history of signals is uncharacteristically positive. Further work by McCardle (1985) and Bhattacharya et al. (1986) has shown that, if information acquisition is costly, then a better new technology might not be adopted: if the early history is uncharacteristically negative, the firm might forgo paying for further information and reject the new technology altogether. Our model differs from these studies because it is strategic. In particular, the "signals" firms receive are incremental innovations that not only provide information, but also affect positions in the race for performance. Both properties are motivated by historical accounts of episodes of disruptive change. They are crucial for our results because they imply that the cost of acquiring further signals (by waiting to adopt) is endogenously determined by competition. Indeed, positions in the performance race determine the propensity to adopt, with leaders more and laggards less prone to adopt. In addition, mistakes in the selection of technologies are driven by market structure and the early history of *relative* performance of the two designs.[6]

Our numerical approach targets Markov Perfect Equilibria in environments where a bounded state vector captures current market structure. This feature is shared with a class of models, surveyed in Doraszelski and Pakes (2007), that follow Ericson and Pakes (1995) and Pakes and McGuire (1994), who analyze industry dynamics with entry and exit. Besanko and Doraszelski (2004) and Doraszelski and Markovich (in press) use a similar framework to study the size distribution of firms and the role of advertising, respectively. These papers are interested in whether asymmetries between firms tend to increase or decrease over time and do not deal with technology choice and learning.

2. The model

2.1. The investment game

We consider a dynamic game played by two players, an incumbent firm (I) and a startup firm (S). Time is discrete, there is an infinite horizon and firms discount the future using the discount factor β. Sales are assumed to depend on *relative product performance*. We represent the performance levels of firms I and S at time t by integers d_t^I and d_t^S, respectively, and denote the incumbent's advantage by $d_t = d_t^I - d_t^S$.

A pair of *revenue functions* $R^k : Z \to \Re_+$; $k = I, S$, maps the incumbent's performance advantage d_t into within-period-profits before investment costs. In Section 2.4 below, revenue functions are derived explicitly from a discrete choice model of Bertrand competition, where d_t indexes the degree of product differentiation. For the time being, we only assume that both functions are nonnegative and bounded, with R^I increasing and R^S decreasing, and that they are symmetric, that is, $R^I(d) = R^S(-d)$ for all d. Figure 1 plots two examples of revenue functions obtained by solving the Bertrand pricing game. Symmetry implies that we only need to plot one revenue function: for each d, the profit of the incumbent can be read off at d and that of the startup at $-d$.

Firms try to improve performance by undertaking investment projects of fixed size. Let $x_t^k = 1$; $k = I, S$ indicate that firm k invests at time t (and $x_t^k = 0$ otherwise). Every project costs c, and can either succeed or fail. The firm's performance level increases by one unit in period $t + 1$, if and only if it undertakes a project in t which then succeeds. Let $z_t^k = 1$ indicate this event. The probability of success of an investment project depends on the *design* the firm

[5] While Jovanovic and Nyarko explicitly model uncertainty about how to operate a new technology, firms in their model are sure that the new technology would be more productive if only they had enough expertise about it. In contrast, the main feature of our model is that *long run* performance growth under the new technology is uncertain.

[6] Existing literature on strategic experimentation (Aghion et al., 1993; Bergemann and Valimaki, 1996; Bolton and Harris, 1999) does not consider technology adoption. These papers are instead concerned with price-setting games, and beliefs are the only state variable. In our model, the dynamics are crucially driven by performance, a payoff-relevant non-belief state variable.

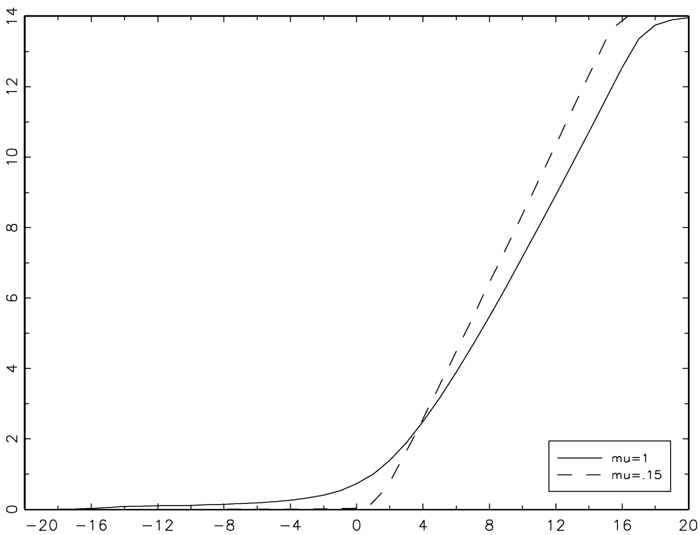

The horizontal axis measures the difference in performance. The vertical axis measures per period profits. The parameter μ governs the substitutability of products of different performance. The other parameters are set to their baseline values: the marginal cost of production is $\kappa = 1$ and the maximum amount a consumer is willing to spend for a unit of the good is $y = 15$. Note that the revenue functions are symmetric: $R^I(d) = R^S(-d)$, so that at each d, the profit of incumbent can be read off at d and that of the startup at $-d$.

Fig. 1. Revenue function of the within-period pricing game.

employs. There are two designs, called *old* and *new*. Each design is identified with a probability of success. The probability p corresponding to the old design is known. The probability q corresponding to the new design is drawn by nature at time zero from a uniform distribution over [0,1] and is unknown to the firms. We denote the realized value of q by q_0. Conditionally on $q = q_0$ and realized action sequences (x_t^I, x_t^S), the random variables z_t^k are independent both over time and across firms.

We assume that the startup firm employs the new design throughout the game. In contrast, the incumbent firm initially employs the old design, but may irreversibly switch to the new design at any time. We denote the adoption choice by $\delta_t \in \{0, 1\}$ and let $\phi_t = 1$ indicate whether adoption has occurred prior to time t. Firms' action spaces are

$$A_t^S = \{0, 1\}, \tag{1}$$

$$A_t^I(\phi_t) := \left\{\left(x_t^I, \delta_t\right) \in \{0, 1\}^2 \colon x_t^I \geqslant \delta_t \text{ and } \delta_t \leqslant 1 - \phi_t\right\} \tag{2}$$

where the first inequality in (2) forces the incumbent to invest when adopting while the second inequality ensures that the incumbent can only adopt once.

We assume that the incumbent's cost of switching from the old to the new design is zero, but that his cost of switching back to the old design are sufficiently large to make the adoption decision irreversible. While these assumptions are extreme, there are reasons to believe that the cost of making a disruptive technological change is lower than the cost of reversing that change. Disruptive changes typically involve major adjustments in a company's organizational capital. For example, new employees who are experts on the new design may have to be hired to replace experts on the old design. The new design will change relationships with suppliers and customers, and it may change also the way different parts of the organization have to cooperate.[7] Changes in organizational capital are costly to reverse if know-how that made up the old organizational capital was destroyed by the disruptive change. Changes in employees and rules of cooperation within the firm and between the firm and its partners may be hard to recover.[8] In contrast,

[7] For example, Henderson (1993) documents communication and coordination costs that arose when innovation in the photolithographic equipment industry changed the way different components of a product, produced by different divisions of a firm, were put together.

[8] Another argument that increases the cost of reversal is based on the incentive of senior managers whose quality is unknown. If managers reverse a major strategic decision, the (less informed) shareholders, suppliers and creditors are likely to revise downwards their estimate about management's ability to plan ahead, which may hurt profits and/or eventually the managers' compensation.

moving to a new design may be less costly if organizational capital can be acquired by copying the startup firm that is already using the new design, for example by hiring away some of its employees.[9]

Players' actions as well as the successes and failures of their investment projects are public information. A typical *history* of play before period t is

$$h^t = \left(x^I_\tau, x^S_\tau, z^I_\tau, z^S_\tau, \delta_\tau; \tau = 1, 2, \ldots, t-1\right).$$

Let H^t denote the set of period t histories. A *strategy* for player k is a collection of functions $\sigma^k = \{\sigma^k_t; \sigma^k_t : H_t \to P(A^k_t(\phi^t));\ t = 0, 1, \ldots\}$; each σ^k_t, maps the histories up to time t into the set of probabilities over actions allowed for player k at time t after that history.

Firms maximize the expected present value of future profits

$$U^I\left(\sigma^I, \sigma^S\right) = E\left[\sum_{t=0}^{\infty} \beta^t \left(R^I(d_t) - cx^I_t\right)\right], \tag{3a}$$

$$U^S\left(\sigma^I, \sigma^S\right) = E\left[\sum_{t=0}^{\infty} \beta^t \left(R^S(d_t) - cx^S_t\right)\right] \tag{3b}$$

where relative performance evolves, for given d_0, according to

$$d_{t+1} = d_t + z^I_t - z^S_t; \quad t \geqslant 1, \tag{4}$$

and expectations are taken under the probability distribution induced by the strategies of the players as well as "nature's strategy" implicit in the distribution of q and the sequences (z^k_t).

2.2. Equilibrium

We are interested in Stationary Markov Perfect Equilibria (SMPEs). The *payoff relevant information* consists of the current performance difference, knowledge of whether the incumbent has adopted, as well as the common belief about the potential of the new design. The latter is summarized by the pair (S, T), where T is the number of investment projects that have been undertaken (by either firm) with the new design, and where S is the number of successful projects. This information is sufficient to form the posterior distribution of q. In particular, the probability of success of an investment project under the new design, conditional on the information (S, T), is equal to the posterior mean $E[q|S,T] = \frac{S+1}{T+2}$ derived from the uniform prior over q. Payoff relevant histories may thus be represented by tuples $\omega_t = (d_t, S_t, T_t, \phi_t)$.

Markov Perfect Equilibria are sequential equilibria such that strategies depend only on payoff relevant histories and time (Maskin and Tirole, 2001). SMPE strategies must also be independent of time. Standard arguments (e.g. along the lines of Pakes and McGuire, 1994) imply that MPEs can be found by solving a pair of Bellman equations. Define the *state space* $\Omega = \{(d, S, T, \phi) \in Z \times Z^2_+ \times \{0, 1\}:\ S \leqslant T\}$.

An equilibrium consists of value functions $V^k : \Omega \to \Re$ and Markov strategies $\sigma^{k*} : \Omega \to P(A^k(\phi))$ such that

1. For $k = I, S$ and with $\tilde\sigma^j = \sigma^{j*}(\omega)$ for $j \neq k$, V^k satisfies

$$V^k(\omega) = \max_{\tilde\sigma^k \in P(A^k(\phi))} \left\{R^k(d) + E\left[-cx^k + \beta V^k(\omega') \mid \omega, \tilde\sigma^k, \tilde\sigma^j\right]\right\} \tag{5}$$

s.t.

$$S' = S + z^S + (\phi + \delta)z^I, \tag{6}$$

$$T' = T + x^S + (\phi + \delta)x^I, \tag{7}$$

$$d' = d + z^I - z^S, \tag{8}$$

$$\phi' = \phi + \delta \tag{9}$$

[9] Costless adoption allows us to isolate the effects of uncertainty on adoption decisions. We have also computed the model with an adoption cost. The qualitative behavior remains similar, with adoption generally taking place later and only after more favorable histories for the new design.

where the actions are distributed according to $\tilde{\sigma}^I$ and $\tilde{\sigma}^S$, i.e. $Prob(x^S = 1) = \tilde{\sigma}^S(\{1\})$, $Prob(x^I = 1) = \tilde{\sigma}^I(\{1,1\}) + \tilde{\sigma}^I(\{1,0\})$, etc., and where the distribution of the project outcomes z^I and z^S *conditional on the information contained in the current state vector* is given by

$$Prob\big(z^I = 1 \mid \omega, \tilde{\sigma}^I, \tilde{\sigma}^S\big) = p\tilde{\sigma}^I\big(\{1,0\}\big) + E[q \mid S, T]\tilde{\sigma}^I\big(\{1,1\}\big), \tag{10}$$

$$Prob\big(z^S = 1 \mid \omega, \tilde{\sigma}^I, \tilde{\sigma}^S\big) = E[q \mid S, T]\tilde{\sigma}^S\big(\{1\}\big). \tag{11}$$

2. σ^{k*} achieves the max in condition 1, for $k = I, S$.

For every state ω, SMPE strategies are a Nash equilibrium of a one shot game with strategy spaces equal to the action spaces for that state and with payoff functions given by the right-hand side of the Bellman equations (5). The value functions in every state are the Nash equilibrium payoffs of this one shot game.

It is straightforward to characterize the equilibrium dynamics in an SMPE. The equilibrium distribution of the random variables z_t^k depends on the strategies chosen by the two players. Since these strategies depend only on the current state ω_t, the equilibrium law of motion for ω_t is a time invariant Markov chain with transition equations given by (6)–(9) and with the distribution of z_t^k induced by σ^{k*}.

2.3. Steady states

Expected profit at any point in time is a weighted average of $R^k(d)$-values, with probabilities determined by the law of motion implied by the strategies. Investment by player k shifts this distribution over $R^k(d)$-values, putting more weight on values more favorable to player k. Since the revenue functions are bounded and increasing, the expected benefit from investing must then be close to zero for values of the performance difference d, that are either high enough or low enough, regardless of what the other player does. As a result, $x^I = x^S = 0$ must be a dominant strategy equilibrium (of the state contingent one-shot game implied by an SMPE) whenever d is large enough in absolute value. Formally, for any SMPE, there exist $d^{hi}$and d^{lo} such that for all $d \geqslant d^{hi}$ or $d \leqslant d^{lo}$ and for all $S, T \in Z_+$, $S \leqslant T$, and $\phi \in \{0,1\}$, $\sigma^{I*}(d, S, T, \phi)$ assigns probability one to the actions $x^I = 0$ and $\delta = 0$ and $\sigma^{S*}(d, S, T, \phi)$ assigns probability one to the action $x^S = 0$.

The set of states ω with $d \geqslant d^{hi}$ or $d \leqslant d^{lo}$ consists of the *absorbing states* of the model. In our numerical examples, we employ a revenue function derived from Bertrand competition with differentiated products. In this case, the set of absorbing states will be reached in finite time with probability one. The model is thus one of an *episode of random length*. After a finite number of periods, a market structure with one dominant firm will be reached.

2.4. Parameterization

We derive revenue functions R^i from a price setting game that is based on the standard logit model of demand (e.g. Anderson et al., 1992). The game justifies our assumption that revenue depends only on *relative* performance, and also shows that relative performance can be identified with market share. The incumbent and startup firm offer versions of an indivisible good that differ in the performance levels d^k. Every period, a continuum of consumers of mass 1 would like to purchase one unit of the good if and only if it is available from some firm at a price below the reservation price y. When consumer m buys a unit of the good from firm k at price p^k, he receives indirect utility $y - p^k + d^k + \epsilon_m^k$. The preference shocks ϵ_m^k are independent across consumers and doubly exponentially distributed with parameter μ, which governs the degree of product differentiation. Firms face equal constant marginal costs κ of producing output.

It is never optimal for either firm to price above y and end up with zero demand. For prices below y, integration over consumers delivers the fraction of consumers who buy from firm k (rather than from $j \neq k$):

$$X^k = \frac{1}{1 + \exp(\frac{d^j - d^k - p^j + p^k}{\mu})}.$$

Demand thus depends only on the difference in performance levels as well as the difference in prices (or, since marginal costs are the same across firms, the difference in markups). The first order condition for firm k's optimal markup is

$$p^k - \kappa = \min\left\{y - \kappa, \frac{\mu}{1 - X^k}\right\}.$$

These conditions characterize Bertrand equilibrium markups in state $d = d^I - d^S$. The revenue functions record the corresponding within-period profits, that is, $R^k(d) = (p_k - \kappa)X^k$. Symmetry of the game implies that $R^I(d) = R^S(-d)$.

The function R^I is plotted in Fig. 1 for two different values of the substitutability parameter μ. The assumption of a maximal markup implies that the profit function is bounded. If the incumbent's performance advantage is large enough, then he can charge the reservation price y and attract almost the entire unit demand; his profit thus tends to the maximal markup $y - \kappa$ as d goes to infinity. At the same time, the follower will never charge less than μ, so that his profit goes to zero as d becomes small. There are two reasons why an increasing and bounded profit function is important for the dynamic game. First, it implies that states where d is large enough in absolute value are stationary states where neither firm invests. Second, it implies that the profit function will be convex for small d and concave for large d. Since this property carries over to the value function, market leaders act in a more risk averse fashion than followers.

The parameter μ governs how steeply the profit function increases for intermediate values of d, and how large profits are when performance is the same. Indeed, if the reservation price is large enough, then profits at $d = 0$ are simply equal to μ. We assume throughout that $\mu < (y - \kappa)/2$, which guarantees that the combined profits of the two firms are increasing in the performance difference between leader and follower. This property is important for the dynamic game because it implies that the steady state is *absorbing*: according to an important unifying principle in models of dynamic competition, identified by Budd et al. (1993), competition tends to evolve in the direction in which the *sum* of payoffs of the competitors is increased. In our setting, constraining revenue function to yield an absorbing steady state is a natural assumption because it makes our model a model of an episode: a market structure with one dominant firm is reached after a finite (but random) number of periods.

Our numerical analysis below has two components. First, we present detailed results for a leading example. For a baseline parameterization, we represent graphically the equilibrium strategies and provide detailed simulation results that condition on various events. The resulting tables and figures make explicit the dynamics of the game. The revenue function for the leading example is the solid line in Fig. 1, with $\mu = 1$, $y = 15$ and $\kappa = 1$. We also assume a discount factor of $\beta = 0.97$, a cost per investment project of $c = 0.5$, and a potential of the old design $p = 0.7$. Second, we demonstrate our main results—the dependence of the adoption decision on market share and the importance of the early history for the persistence of market power—for an array of parameterizations. The resulting comparative statics show how the effects we stress vary with the different parameters for the results. They also illustrate that the main qualitative patterns are similar across many different versions of the model. We have also verified that as we move in small steps from the baseline parameterization to the other parameterizations reported below, the equilibrium decision rules change continuously.

3. Characterizing equilibria

This section describes equilibrium strategies and payoffs in two steps. Section 3.1 considers the subgame that begins when the incumbent adopts the new design. In other words, it describes the region of the state space where $\phi = 1$. Section 3.2. then characterizes the pre-adoption phase, where the two designs compete head-on ($\phi = 0$).

3.1. Racing with equal designs after adoption

After adoption, firms enter a race in which they invest to generate incremental innovations, both operating under the new design. Since $R^I(d) = R^S(-d)$, this one-technology race is symmetric—formally, we have $V^I(d, S, T, 1) = V^S(-d, S, T, 1)$). States d and $-d$ thus lead to identical strategy combinations up to the identity of the players. Firms trade off the costs of investment against both the gain from improved performance and the value of the additional information to be gained by investing. To isolate the "performance race" dynamics from the learning dynamics, we first report results for the case where the probability of success is known.

3.1.1. The certainty-one-technology-game

Figure 2 shows how equilibrium strategies depend on d, together with the value function, for a known probability of success. There are three equilibrium regions. As argued in Section 2.3, for d high enough both firms will not invest. Moreover, both firms invest at zero—this always happens if the cost of investment is not too high. Finally,

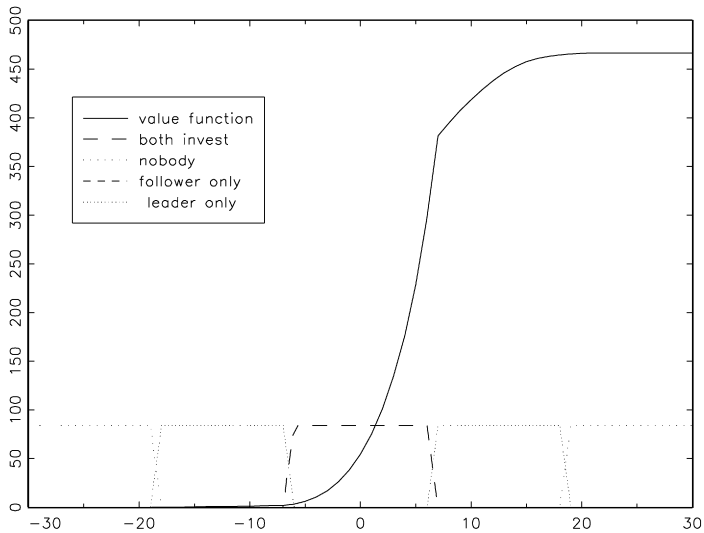

The horizontal axis measures the difference in performance; the vertical axis measures the firm's corresponding value function. The different broken lines indicate equilibrium action regions. Parameter values: probability of success $q_0 = 0.7$, cost of investment in performance accumulation $c = 0.5$, $\kappa = 1$, $y = 15$, $\mu = 1$.

Fig. 2. Equilibrium value function and equilibrium regions for the one technology certainty game.

at intermediate values of d the firm that has fallen behind in performance stops investing, while the leader keeps increasing his advantage.

This structure follows from the curvature properties of the revenue function. Figure 1 shows that the revenue function is steeper at a typical state $d > 0$ than at the negative counterpart of that state, $-d$. In particular, there exist values of $d > 0$ such that (i) the revenue function is steep enough at d so that the leader keeps investing to increase market share further even if the follower does not invest, while (ii) the revenue function is flat enough at $-d$ so that the follower has no incentive to invest, even if the leader does still invest. The model thus predicts that the investment intensity of followers is never higher than that of leaders.

The equilibrium dynamics of performance (and hence market shares) can be read directly off Fig. 2. If d_t is in the region where both firms invest, it evolves as a random walk without drift. During this "Schumpeterian" phase of the race, both firms improve performance and this growth is accompanied by frequent changes in market shares. Eventually, d_t hits one of the two boundaries and enters an "innovative leadership" phase. Here only the market leader invests, increasing his market share further, while the follower's product loses ground—d_t continues moving away from zero with a drift, until it finally stops once it hits the boundary of the no-investment region. A steady state is thus reached in finite time with probability one. This pattern is common to all our examples: it is what makes the model one of an episode of technological change (with random end). We refer to the long run leader as the "winner" of the race.

3.1.2. *Growth and variance effects of changing potential*

To understand the impact of learning about the potential q_0 of a technology, it is helpful to first consider how the investment dynamics depend on potential when the latter is known.[10] Equilibrium actions for certainty-one-technology-games with different q_0 are shown in the bottom panel of Fig. 3, with the corresponding value functions in the top one. The performance advantage of one player, say the incumbent, is measured along the horizontal axis and every row of rectangles corresponds to a different value of q_0. Every individual rectangle indicates the equilibrium actions taken in one state of the game, where the top-right half of the rectangle represents the action of the incumbent,

[10] This is also of interest in itself: potential here measures the average number of successful incremental innovations per investment project—a measure of R&D productivity. We can therefore analyze how R&D productivity shapes market structure in our setting.

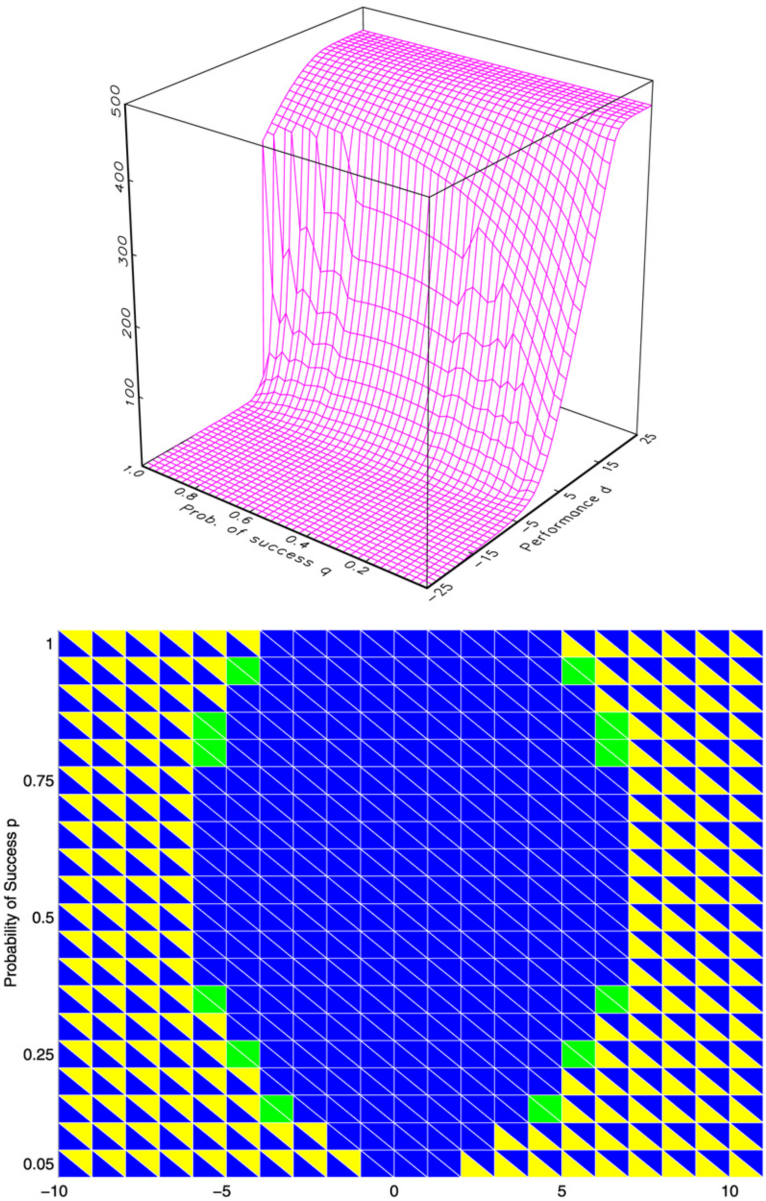

The top panel plots value functions for "certainty, one technology" games that differ by the probability of success under the one, known, design. The only argument of these value functions is the performance advantage.
The bottom panel illustrates equilibrium actions in "certainty, one technology" games that differ by the probability of success. The vertical axis measures the probability of success; every row of squares thus describes a different game. The horizontal axis measures the performance advantage of the incumbent, the single state variable of the game. Every individual rectangle summarizes the actions in one state. The bottom-left triangle indicates the action of the startup, the top-right triangle that of the incumbent. Actions are color coded—dark (blue in the web version) for investment and light (yellow in the web version) for no investment. Gray (green in the web version) rectangles indicate mixed actions.

Fig. 3. Certainty-one-technology-game: comparative statics.

and the bottom-left half that of the startup. In both cases, a dark (blue in the web version) triangle indicates investment, while a light (yellow in the web version) triangle indicates no investment.

The investment region is widest for $q_0 = \frac{1}{2}$ and narrows as q_0 moves toward zero or one, where the narrowing is more pronounced for low q_0. To understand this pattern, consider the dual impact of changing the potential parameter q_0. First, there is a *growth* effect, since q_0 governs the average growth rate of performance. An episode in which performance *levels* grow faster on average corresponds to a higher value of q_0. Second, there is a *variance* effect, since q_0 determines the random walk behavior of d_t. Indeed, when both firms invest, the shocks driving d_t have mean zero and variance $2q_0(1-q_0)$. The variance of shocks is thus hump-shaped in q_0, with a maximum at $q_0 = \frac{1}{2}$.

The growth effect implies that the leader's expected profits during the "innovative leadership" phase are increasing in q_0, because it is on average less costly to further increase the performance advantage. If this effect alone was at work, we would expect to see the investment region widening for higher q_0. The hump shape is induced by the variance effect: during the "Schumpeterian" phase, intermediate values of potential (that is, q_0 close to 1/2) induce a greater variance of the d_t process and therefore more turbulence in market shares. This increases the follower's investment intensity, because for low values of d the value function is convex; the follower exhibits risk-loving behavior. He will thus keep investing longer (i.e. even if he is more behind) the higher the variance of shocks to the performance process.

3.1.3. The role of learning

If the potential of the new design is uncertain, investment decisions must take the gradual revelation of information into account. Figure 4 compares equilibrium regions with and without uncertainty about potential. The right-hand set of panels shows the equilibrium strategies when firms learn about q, as a function of the three state variables d, S and T. This is again a symmetric game. The horizontal axis measures the performance advantage of the race leader.

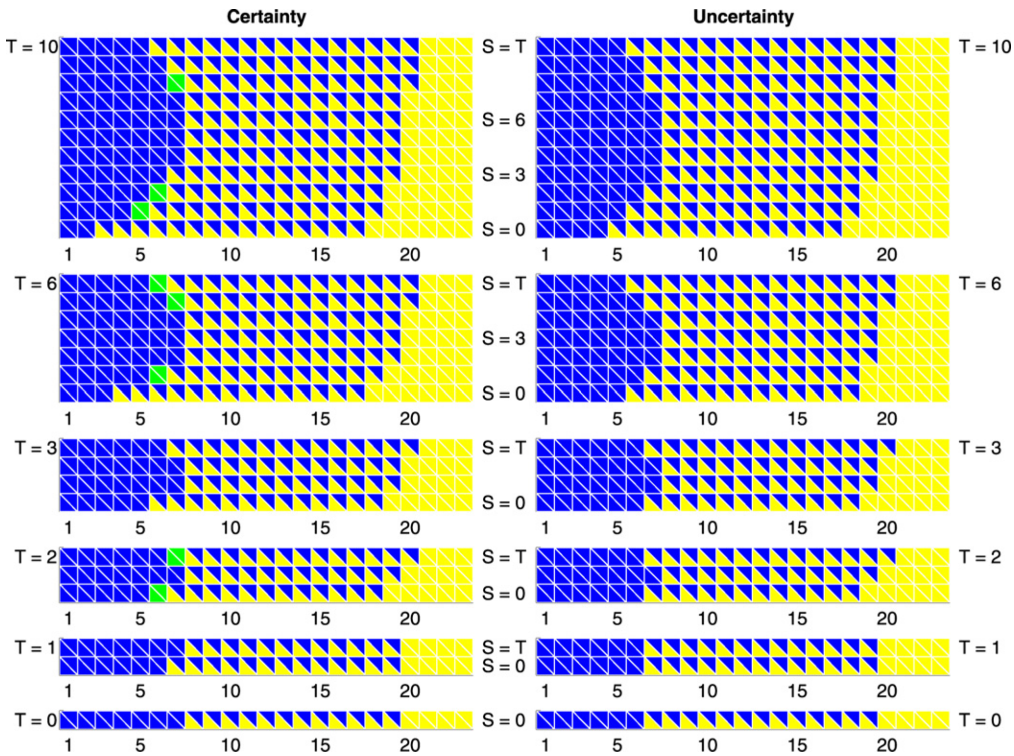

The right-hand set of panels describes equilibrium actions in the uncertainty-one-technology game as a function of the state variables: the performance advantage of the race leader and the beliefs about the new design. The horizontal axis measures the performance advantage of the leader. Each individual panel is identified by a number T of trials observed for the new design, and the vertical axis counts the number S of successes in these trials. Every line thus corresponds to a different belief state.

For comparison, the left-hand panels report equilibrium actions in a set of certainty-one-technology games. Every line describes one such comparison game, with its (known) probability of success chosen to be equal to the posterior mean $\frac{S+1}{T+2}$ in the corresponding line (that is, belief state) of the uncertainty game.

In all panels, an individual rectangle describes actions in one state of the game. The bottom-left triangle indicates the action of the follower, the top-right triangle that of the leader. Actions are color coded—dark (blue in the web version) for investment in the old technology and light (yellow in the web version) for no investment. Gray (green in the web version) rectangles indicate mixed actions.

Fig. 4. Equilibrium structure: certainty and uncertainty one-technology games.

As before, every individual rectangle indicates equilibrium play in one state of the game, where the top right triangle indicates the action of the race leader, and the bottom left triangle that of the follower.

Every individual row corresponds to one "belief state" (S, T). Every individual panel corresponds to a value of T—the number of investment projects undertaken using the new design, and therefore also the number of signals observed about the new design. Every row of rectangles within a panel corresponds to a value of S, the number of successful investment projects under the new design. Rows higher up within a panel thus correspond to states where firms are more optimistic about the potential of the new design. For comparison, the left hand panel presents equilibria for a set of comparison games in which q is known. For every row corresponding to a belief state (S, T) we have set the probability q_0 equal to the expected potential (the posterior mean of q), $\frac{S+1}{T+2}$ in that belief state.

The basic pattern observed in Fig. 3 is also present in the uncertainty case: the investment regions narrow as the posterior mean moves away from one half. What changes with uncertainty now depends on the posterior mean of q: if the posterior mean is greater than $\frac{1}{2}$, the investment region is *wider* than in the certainty case, whereas the opposite effect obtains for a posterior mean below half. This result can be understood in light of the growth and variance effects discussed above.

Indeed, followers are willing to fight harder (that is, invest even when they have a substantial performance disadvantage) if *either* the expected growth rate in the leadership phase is higher (higher success probability), *or* if the variance of d_t is higher. Now if the posterior mean is low, then there is some chance that both the growth rate and the variance are high. The follower thus fights harder than in the certainty case where both are known for sure. In contrast, if the posterior mean is high, then there is some chance that the variance is higher, but there is also a chance that the expected growth rate is lower. The latter effect dominates, which makes the follower fight less hard.

3.2. Racing with different designs and the adoption decision

Figure 5 shows equilibrium actions in the pre-adoption states ($\phi = 0$). The structure of this figure is the same as that of Fig. 4—the right-hand panels show the game with unknown potential of the new design q, while the left-hand panels report actions in a set of comparison games with known q. In the comparison game for a belief state (S, T), q is equal to the posterior mean $\frac{S+1}{T+2}$ in that belief state. One difference from Fig. 4 is that the game is no longer symmetric for $\phi = 0$; we thus plot equilibria for both positive and negative values of d. Moreover, individual triangles now indicate the design firms employ when investing: dark (blue in the web version) for the old design, and gray (green in the web version) for the new design. White triangles indicate no investment.

3.2.1. Immediate technology choice under certainty

Consider first the adoption decision when the potential of the new design q_0 is known. The left-hand panels of Fig. 5 show that the incumbent *immediately* adopts the new design if and only if its potential is higher than that of the old technology ($q_0 > p$). For example, consider the left-hand panel showing equilibria for $T = 6$. Its seven rows show what happens for $q_0 = (S + 1)/8$, for $S = 0, 1, \ldots, 6$. Only the top two rows (that is $q_0 = 7/8$ and $q_0 = 6/8 = 0.75$) contain equilibria where the incumbent employs the new technology: these rows are the only ones where $q_0 > p = 0.7$.

It follows that a model without uncertainty cannot capture design competition, where both firms simultaneously employ, and invest in, different technologies. For $q_0 \geqslant p$, the equilibrium dynamics are exactly the same as in the one-technology race described above. For $q_0 < p$, the dynamics is different since the process d_t now has a positive drift as long as both firms invest. This increases the incumbent's probability of winning.

3.2.2. The propensity to adopt while learning

The adoption decision in the uncertainty game can be summarized by two rules-of-thumb. First, it is natural that the incumbent adopts the new design if its expected potential is high enough: in the right-hand panels of Fig. 5, the gray (green in the web version) upper triangles that indicate adoption by the incumbent are located in the top rows that correspond to high expected potential. Second, the incumbent adopts the new design if his performance advantage is large enough: in a given row, the gray (green) triangles tend to be to the right, where the incumbent advantage d is higher.

This second property is less obvious: why should an incumbent who is further ahead in terms of performance have a higher propensity to adopt the new design? The answer follows from the incumbent's attitude towards risk. Consider the tradeoff he faces in any given state in the investment region: he can either choose to keep playing the game he is

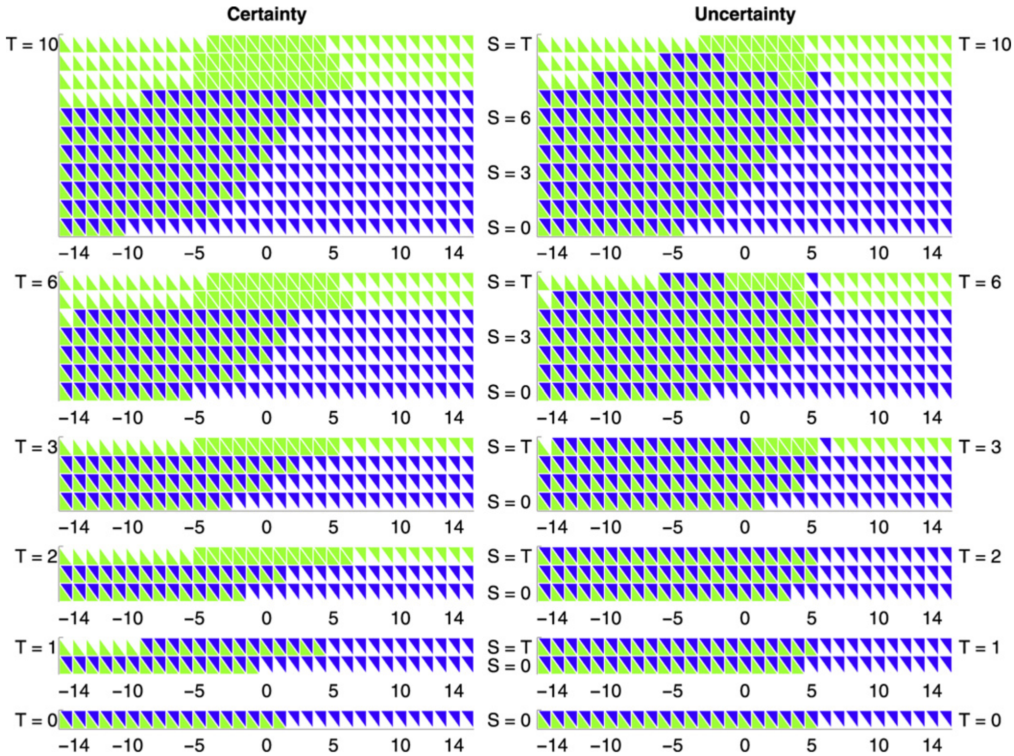

The right-hand set of panels describes equilibrium actions in the full model, with uncertainty about the new design, as a function of the state variables: the performance advantage of the incumbent and the beliefs about the new design. The horizontal axis measures the performance advantage of the incumbent, $d^I - d^S$. Each individual panel is identified by a number T of trials observed for the new design, and the vertical axis counts the number S of successes in these trials. Every line thus corresponds to a different belief state.

The left-hand set of panels reports equilibrium actions in a set of comparison games in which the probability of success under the new design is known. Every line describes one such game, with its probability of success set equal to the posterior mean $\frac{S+1}{T+2}$ in the corresponding line (that is, belief state) of the full model.

In all panels, every individual rectangle corresponds to one state of the game. The bottom-left triangle indicates the action of the startup, the top-right triangle that of the incumbent. Actions are color coded—light (green in the web version) for investment in the new technology, dark (blue in the web version) for investment in the old technology, and white for no investment. In both games the probability of success of the old technology p is set to 0.7.

Fig. 5. Equilibrium structure: full model.

in, which features an unknown drift in d_t which could work for or against him, or he can opt for a game with zero drift.

Two effects make the old design attractive for low values of d. First, the convexity of the revenue function for low values of d induces risk-loving behavior, here on the part of an incumbent. At low values of d, the incumbent thus prefers the riskier strategy of staying with the old design (he essentially "gambles for resurrection"), whereas for higher values of d the concavity of the revenue function makes him risk averse, and thus more likely to adopt. Second, the irreversibility of the adoption decision induces an *option value* of staying with the old design, formally defined as the difference between the value of staying with the old design in the present game minus the value of the same action in a hypothetical game in which the incumbent is allowed to freely switch back and forth between the old and new technologies. This option value is higher for low values of d: staying with the old design keeps alive the possibility of getting a drift that works in favor of the player, which is more important for a player with a low (or negative) advantage.

3.2.3. Adoption and expected potential

To highlight the role of risk considerations for the adoption decision, we compare the adoption rule in the model to a simple "expected potential" rule (EPR) that says: employ the new design if and only if its expected potential is higher than the known potential of the old design. As discussed above, the EPR is optimal when the potential of the new design is known. We now show that, under uncertainty, there are deviations from the expected potential in both directions: incumbents may either *not* adopt a new design that is expected to be better than the old design, or they may adopt a new design that is expected to be worse than the old design.

The first type of deviation is already present in the baseline case shown in Fig. 5. For example, consider the equilibrium regions in the right-hand panel corresponding to $T = 6$. In the first two rows, the expected potential of the new technology is $E[q] = 0.875$ and $E[q] = 0.75$, respectively, strictly larger than the potential of the old design, $p = 0.7$. Nevertheless, there are many states in the first two rows where the incumbent does not adopt the new design. For incumbents who have fallen behind in the performance race, the benefit of a more risky strategy and the option value of waiting outweigh the losses in expected performance growth.

We have numerically explored deviations from the EPR for a large set of alternative parameterizations. In our calculations, we vary the substitutability parameter μ in the interval $[\frac{1}{3}, 3]$ the cost of investment $c \in [0.1, 2.5]$ and the discount factor $\beta \in [0.87, 0.98]$. We compare models on behavior in two sets of states. First, we consider "EPR-wait-states" where the new design is expected to be worse than the old design, so that the EPR recommends to wait rather than adopt. In particular, we select states with $(S, T) = (1, 2)$, $(2, 4)$ or $(3, 6)$, where the posterior mean is $E[q] = 0.5$, and we set the potential of the old design to $p = 0.51$. We have experimented with larger values of p, but deviations from the EPR in EPR-wait-states were never found for values of p more than a few percentage points larger than one half. Second, we consider "EPR-adopt-states" where the new design is expected to be better than the old design ($p = 0.5$), so that the EPR recommends adoption. In particular, we consider states with $(S, T) = (2, 2)$, $(3, 4)$, $(4, 6)$ or $(5, 6)$.

Comparison of behavior in the two sets of states delivers three basic patterns. First, for *all* parameterizations, there are deviations from the EPR in at least some of the EPR-adopt-states. Non-adoption of a new design, even if its expected potential is higher than the old one, is thus a robust feature of the model. Second, there exist EPR-wait-states in which the optimal decision is actually adoption. In other words, it may be optimal to adopt a new design that is expected to be worse than the old one. Third, decreasing β or μ from the upper bound, to the baseline value, and then to the lower bound (i) increases the investment region, (that is, the region of d values where both firms invest, holding fixed S, T)) and (ii) reduces the number of pairs (S, T) for which deviations from the EPR occur. Results (i) and (ii) also follow if c is increased. This pattern says that both types of deviations from the EPR tend to occur in games where competition is "fiercer." Indeed, wider investment regions imply that both firms continue to invest even if one already has a large performance advantage.

Intuitively, risk considerations are more important when the game is expected to last longer. In fact, an incumbent with $p > E[q]$ is more reluctant to give up the old design if he knows that the startup might soon stop investing, because the risk of catching up is reduced; instead, if he expects the startup to keep investing for a long time, he will be more reluctant to run the risk of racing with different designs. It is also intuitive that we observe more deviations from the EPR in EPR-adopt-states. Indeed, there are two forces pulling the incumbent away from the EPR which need not work in the same direction. The first force is that adoption reduces risk, which leads to more adoption for low d when the incumbent is risk-loving, but to less adoption for high d when the incumbent is risk averse. The second force is that there is an option value of waiting, which *always* leads to less adoption. Both forces work together to generate non-adoption of a better design when the incumbent has fallen behind. This leads to many deviations in EPR-adopt-states. In contrast, the two forces work in opposite directions when the incumbent is ahead. To obtain a deviation in an EPR-wait-state, risk must matter enough so that the first force dominates and more than wipes out the option value.

4. Simulation results

The simulation results of this section show how the steady state outcomes of an episode, that is adoption and persistence (or loss) of market power depend on initial conditions, parameter values and the early history of performance under the new design. The main point is that changes in market leadership—i.e. episodes in which the startup eventually wins and replaces the incumbent as the market leader—are often associated with a weak early history of the new design. This point is established in Section 4.1. Section 4.2. considers the role of other factors on the startup's probability of winning and Section 4.3 deals with technology selection.

Throughout this section, we view the model from the perspective of an observer who compares whole episodes, that is, full equilibrium sample paths produced by the model. A set of summary statistics for the baseline parameterization is reported in Table 1. The panels and broad columns in the table correspond to different values of the potential of the

Table 1
Simulation results: summary statistics

d_0	$p = 0.5$																				
	$q_0 = 0.325$			$q_0 = 0.5$			$q_0 = 0.575$			$q_0 = 0.7$			$q_0 = 0.775$			$q_0 = 0.85$			$q_0 = 0.925$		
0	0.95	0.9	0	0.56	0.41	0.1	0.41	0.18	0.08	0.34	0.01	0.01	0.35	0.01	0.01	0.33	0	0	0.28	0.01	0.01
	1	1	0	0.51	1	0.49	0.52	0	0	0.48	0	0	0.5	0	0	0.48	0	0	0.5	0	0
1	0.94	0.85	0	0.64	0.38	0.05	0.53	0.16	0.06	0.47	0.02	0.01	0.47	0.01	0	0.47	0	0	0.46	0	0
	1	1	0	0.56	1	0.44	0.59	0	0	0.59	0	0	0.63	0	0	0.6	0	0	0.63	0	0
2	0.9	0.71	0	0.65	0.28	0.02	0.6	0.15	0.03	0.59	0.03	0	0.57	0	0	0.61	0	0	0.66	0	0
	1	1	0	0.62	1	0.38	0.64	0	0	0.65	0	0	0.67	0	0	0.65	0	0	0.78	0	0
3	0.85	0.62	0	0.74	0.29	0.01	0.69	0.14	0.01	0.69	0.03	0	0.67	0.01	0	0.72	0	0	0.79	0	0
	1	1	0	0.71	1	0.29	0.73	0	0	0.69	0	0	0.73	0	0	0.77	0	0	0.89	0	0
4	0.9	0.58	0	0.81	0.27	0	0.78	0.19	0.01	0.74	0.04	0	0.82	0.02	0	0.84	0	0	0.91	0	0
	1	1	0	0.81	1	0.19	0.77	0	0	0.81	0	0	0.85	0	0	0.88	0	0	1	0	0
5	0.94	0.61	0	0.9	0.44	0	0.87	0.3	0	0.88	0.18	0	0.91	0.16	0	0.93	0.06	0	0.95	0.04	0
	1	1	0	0.83	1	0.17	0.86	0	0	0.88	0	0	0.92	0	0	0.96	0	0	1	0	0
6	1	1	0	1	1	0	1	1	0	1	1	0	1	1	0	1	1	0	1	1	0
	1	1	0	0.94	1	0.06	0.93	0	0	0.92	0	0	1	0	0	1	0	0	1	0	0
d_0	$p = 0.7$																				
	$q_0 = 0.325$			$q_0 = 0.5$			$q_0 = 0.575$			$q_0 = 0.7$			$q_0 = 0.775$			$q_0 = 0.85$			$q_0 = 0.925$		
0	1	1	0	0.98	0.96	0	0.93	0.9	0.01	0.56	0.51	0.12	0.37	0.27	0.18	0.27	0.21	0.2	0.23	0.17	0.17
	1	1	0	1	1	0	1	1	0	0.48	1	0.52	0.5	0	0	0.51	0	0	0.5	0	0
1	0.99	0.99	0	0.96	0.92	0	0.91	0.82	0.01	0.63	0.46	0.07	0.47	0.17	0.08	0.41	0.07	0.06	0.42	0.07	0.07
	1	1	0	1	1	0	1	1	0	0.57	1	0.43	0.59	0	0	0.6	0	0	0.66	0	0
2	0.99	0.98	0	0.96	0.88	0	0.91	0.77	0	0.72	0.41	0.03	0.59	0.17	0.04	0.58	0.05	0.01	0.6	0.02	0.01
	1	1	0	1	1	0	1	1	0	0.63	1	0.37	0.63	0	0	0.67	0	0	0.75	0	0
3	0.99	0.97	0	0.94	0.83	0	0.88	0.73	0	0.78	0.39	0.02	0.69	0.17	0.01	0.72	0.06	0	0.78	0.02	0
	1	1	0	1	1	0	1	1	0	0.69	1	0.31	0.78	0	0	0.78	0	0	0.87	0	0
4	0.99	0.96	0	0.94	0.83	0	0.94	0.77	0	0.87	0.49	0.01	0.84	0.33	0	0.83	0.17	0	0.84	0.06	0
	1	1	0	1	1	0	1	1	0	0.78	1	0.22	0.84	0	0	0.86	0	0	1	0	0
5	0.99	0.95	0	0.98	0.87	0	0.97	0.78	0	0.91	0.58	0	0.91	0.43	0	0.92	0.33	0	0.96	0.16	0
	1	1	0	1	1	0	1	1	0	0.87	1	0.13	0.91	0	0	0.96	0	0	1	0	0
6	1	1	0	1	1	0	1	1	0	1	1	0	1	1	0	1	1	0	1	1	0
	1	1	0	1	1	0	1	1	0	0.93	1	0.07	1	0	0	1	0	0	1	0	0

The first sub-column reports the percentage of cases in which the incumbent “wins,” the second the percentage in which the incumbent does not adopt the new technology and the third the percentage in which it is bypassed, i.e. does not adopt and is displaced by the startup. The first row reports the results for the uncertainty game, while the second for the certainty one. Each run is based on a sample of 500 cases.

old design p and the potential of the new design q_0, respectively. Pairs of rows correspond to values of the incumbent's advantage d_0. For every cell p, q_0 and d_0, the first row reports three probabilities for the outcome of the full game with learning: (i) the probability that the incumbent wins, (ii) the probability that the incumbent does not adopt, and (iii) the probability that the incumbent does not adopt and loses. For every p, q_0 and d_0, the second subrow reports the same numbers (i)–(iii) for the corresponding certainty game.

4.1. *The early history and changes in market power*

When is it more likely that the incumbent retains market dominance? The answer follows from comparing entries in the first subrow and first subcolumn in Table 1 across rows that correspond to different values of d_0 (but holding fixed the broad column and thus q_0). Naturally, the incumbent is more likely to win (that is, to remain market leader at the end of the episode) if he starts out from a higher initial advantage d_0. At the same time, the probability that the incumbent wins as a function of the potential q_0 of the new design is U-shaped. This pattern is somewhat surprising: one would expect that a better new design favors the startup. We now show how this result is driven by learning: a very good new design will quickly reveal its potential to the incumbent, who thus adopts earlier, before dissipating his initial advantage, and increases his chances of retaining market leadership.

4.1.1. *The role of the early history*

Learning implies that the adoption decision of the incumbent depends not only on the true potential of the new design, but also on the *early history* of trials. In Fig. 6, we relate the final market structure to both factors. In both panels, the horizontal axis measures the potential of the new design, while the vertical axis measures the difference between the average fraction of successes S/T between episodes where the startup wins and episodes where the incumbent wins. A point above zero on the vertical axis thus indicates that the early history of the new design is better when the startup wins than when the incumbent wins. In contrast, a negative value means that startup wins tend to be associated with worse early histories for the new design than incumbent wins. Calculations are based on our baseline parameterization with an initial performance difference $d_0 = 3$. We also focus on the pre-adoption phase: the top and bottom panels consider periods 1 and 3, respectively: under the given parameterization, the incumbent never adopts before period 4. The difference in the fraction of successes is reported for two sets of paths: the dashed line takes into account all paths, while the solid line conditions on paths where both players are still investing in the period under consideration and where the incumbent eventually adopts the new technology.

The solid lines show that startup wins that follow adoption by the incumbent are associated with worse early histories for the new design. Table 1 shows that cases in which the incumbent ends up adopting make up 83% of all paths for $q_0 = 0.775$ and more than 94% for $q_0 \geqslant 0.85$. The dashed lines average also across episodes where the incumbent does not adopt. For sufficiently good new designs ($q_0 \geqslant 0.88$), startup wins are associated with worse early histories for the new design (whether or not the incumbent adopts), whereas they are associated with better histories for the new design for lower values of q_0. Table 1 shows that the fraction of startup wins when there is no adoption (recorded in the third subcolumn of each broad column) is close to zero. The difference between the two lines for low q_0 is thus due to paths along which the incumbent does not adopt *and wins*. On such paths, the startup tends to do so badly that he quits investing before the higher potential of his design has been revealed. Going from the solid to the dashed line adds paths without adoption, and thus lowers the fraction of startup successes along paths where the incumbent wins. Startup wins can then be associated with relatively better early histories. The difference between the two lines diminishes as q_0 increases and the probability of adoption goes to one.

We have seen that, unless the early history is so bad that the incumbent wins without adopting, a bad early history favors the startup. The reason is that a bad early history delays the adoption decision. This is illustrated in Fig. 7. The dashed (solid) line reports the average adoption time for paths along which the startup (incumbent) eventually wins, as a function of the potential of the new design. Both lines condition on the incumbent adopting the new technology before the game stops. Along paths where the startup wins, the incumbent adopts on average about 2 periods later. Table 2 relates the adoption time and also the incumbent's advantage at the time of adoption to the true potential of the new design. Consider the cases with $p = 0.5$: the advantage retained by the incumbent at the time of adoption is larger if the new design has a higher probability of success, as this is more likely to induce an early adoption. This is remarkable, since a higher q_0 also implies faster performance accumulation for the startup before adoption.

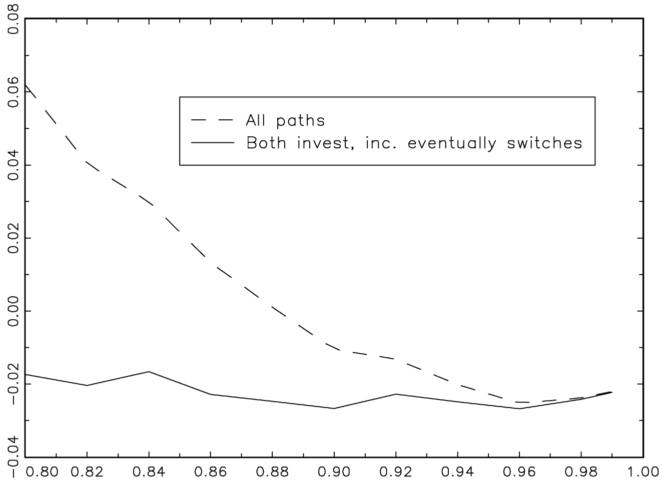

(a) Period 1

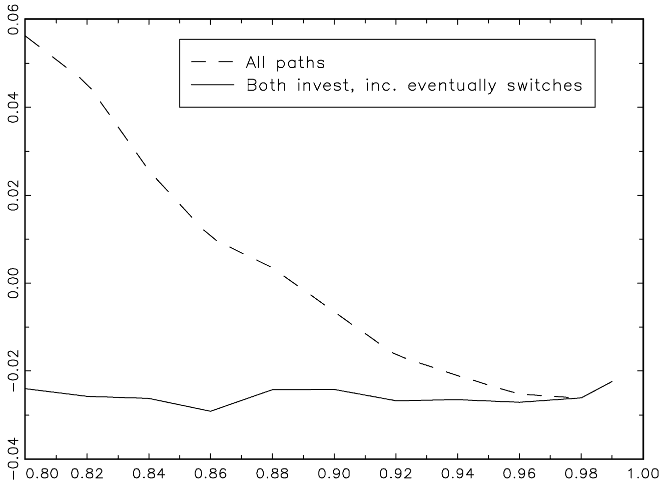

(b) Period 3

The horizontal axis measures different values of q_0; the vertical axis measures the difference in performance (average No. of successes) of the new technology between cases in which the startup eventually wins and loses. The upper panel report the difference after 1 period, the lower after 3 periods. For this game, no adoption occurs before period 4, so all the cases reported are pre-adoption outcomes. Parameter values: $p = 0.7$, $d_0 = 3$, cost of investment in performance accumulation $c = 0.5$, $\kappa = 1$, $y = 15$, $\mu = 1$.

Fig. 6. Differences in new technology performance for paths on which the startup wins and loses.

Table 2
Average adoption time and advantage at the time of adoption, baseline case

d_0	$p = 0.5$				$p = 0.7$	
	$q_0 = 0.575$		$q_0 = 0.85$		$q_0 = 0.85$	
0	16.19	−1.66	5.38	−1.52	15.07	−1.51
1	11.15	−0.34	4.57	−0.29	12.16	−0.57
2	8.31	0.98	3.40	1.22	8.54	0.75
3	5.52	2.30	2.62	2.39	7.15	1.94
4	4.07	3.27	2.41	3.41	5.76	2.99
5	2.78	4.36	2.20	4.45	5.09	3.98

The first column reports the average number of periods before adoption and the second the average performance advantage at the time of adoption for the cases in which adoption takes place.

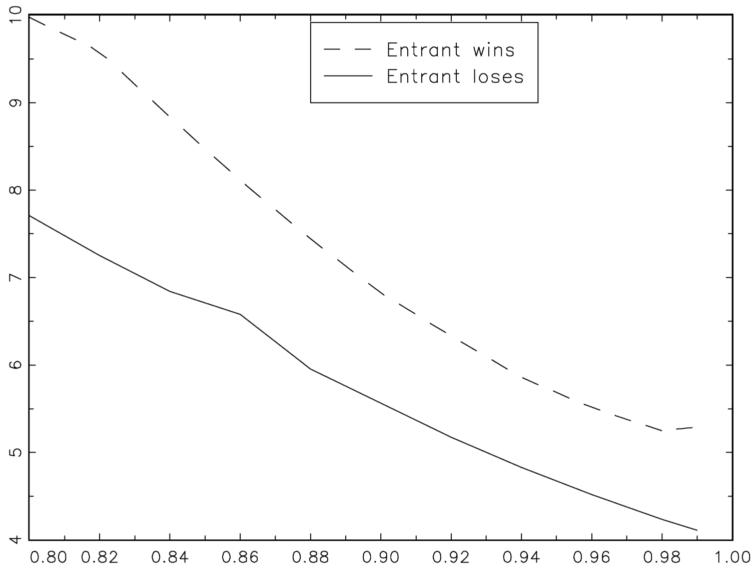

The horizontal axis measures different values of q_0; the vertical axis measures the incumbent's average adoption time on paths in which the startup eventually wins and loses. Parameter values are as in Fig. 6.

Fig. 7. Delay in adoption when the startup wins and loses.

Putting these effects together, we can now explain why the incumbent's probability of winning is a U-shaped function of q_0. On the one hand, a very bad initial performance can result in the startup stopping investing and the incumbent retaining market dominance without adopting. On the other hand, a very good initial performance signals that the new technology is superior and induces early adoption. Chances are highest for the startup when the new design does well enough to prevent an early stop on his side, but not too well to induce early adoption. These conditions are more likely to be met for intermediate values of q_0.

4.1.2. *Alternative parameterizations*

Table 3 illustrates how the role of the early history depends on the parameters of the game. Each row displays a game where one parameter has been changed with respect to the baseline parameterization. As in Fig. 6, panel (b), we report the difference, as of period 3, in the fraction of successes under the new design between paths along which the startup wins and paths along which the incumbent wins.[11] A negative number indicates that the performance of the new technology (the early history) is worse when the startup eventually wins. We report results both for all paths and for paths where both firms invest in period 3 and the incumbent eventually adopts, but has not adopted by period 3.[12]

The results confirm the basic patterns from the baseline case. First, for a new design that is better than the old one and that the incumbent eventually adopts, a startup win is associated with a worse early history than an incumbent win. Second, for sufficiently good new designs, startup wins are associated with worse early histories, whether or not the incumbent adopts. The difference between the two lines arises because a bad early history is useful only as a signal that delays adoption; paths where the incumbent does not adopt tend to be associated with an incumbent win following an early stop by the startup, due to a bad initial performance.

When parameters change in a way to make competition less fierce (in the sense of narrowing the investment regions and hence shortening the expected length of the game), the early histories associated with (unconditional) startup wins become better. Indeed, as the initial advantage d_0 or the cost of investment c increase, or as the substitutability parameter μ or the maximal price y decrease, the unconditional performance difference (the first row of numbers in Table 3) increases. At the same time, when parameters change to make competition less fierce, then the early histories

[11] Results are similar when considering period 1 or 2.

[12] The last qualifier was redundant in Fig. 6, as adoption never occurred before period 3. In some of the cases of Table 3, such as for $p = 0.5$, the condition is binding.

Table 3
Initial history: other parameterizations

	Values of q_0										
	0.5	0.55	0.6	0.65	0.7	0.75	0.8	0.85	0.9	0.95	0.99
Baseline					17.7	10.9	5.7	1.7	−0.6	−2.3	−2.2
					−1.8	−2.4	−2.7	−2.7	−2.4	−2.7	−2.2
$d_0 = 1$					12.2	6.3	2.2	0.0	−0.9	−0.6	0.0
					−0.7	−0.8	−1.0	−1.2	−1.0	−0.5	0.0
$d_0 = 5$					39.1	31.7	25.1	17.6	10.1	3.1	−1.2
					−2.9	−3.8	−3.8	−4.1	−3.9	−3.8	−2.5
$c = 0.25$					9.8	4.1	1.1	0.1	−0.5	−0.9	−0.7
					−0.4	−0.7	−1.0	−0.8	−0.8	−0.9	−0.7
$c = 1$					33.8	27.2	20.4	14.4	7.9	2.3	−0.7
					−1.6	−2.2	−2.7	−2.7	−2.9	−2.6	−1.6
$\mu = 0.5$					19.0	12.2	6.7	2.4	−0.7	−2.6	−1.7
					−1.4	−2.0	−2.2	−2.5	−2.6	−2.7	−1.6
$\mu = 2$					8.9	2.2	0.3	−0.1	−0.4	−0.5	−0.2
					0.0	−0.6	−0.5	−0.4	−0.3	−0.1	0.1
$y = 7.5$					18.7	11.6	6.4	2.5	−0.2	−1.6	−2.2
					−0.9	−1.7	−2.2	−2.5	−2.3	−2.1	−2.2
$y = 30$					17.1	9.8	4.7	1.1	−1.0	−2.4	−2.3
					−1.4	−1.9	−2.5	−2.6	−2.4	−2.7	−2.2
$p = 0.8$						9.8	4.7	1.1	−1.0	−2.4	−2.3
						−1.9	−2.5	−2.6	−2.4	−2.7	−2.2
$p = 0.5$	10.3	6.0	3.4	1.3	0.6	−0.9	−1.7	−2.3	−2.5	−2.8	−3.6
	−1.8	−2.2	−1.9	−2.2	−1.4	−1.7	−0.8	−1.5	−1.4	−0.8	1.1

The table reports the differences in new technology performance for paths on which the startup wins and loses for different parameter configurations for cases with $q_0 \geqslant p$. We report the difference for "All paths" (first line) and for "Both invest, inc. eventually switches" (second line). The baseline reported in the first two lines is that of Fig. 6. We change one parameter at a time; the modified parameter value is reported in the first column.

associated with startup wins after adoption become worse: the second row of numbers shifts down. Intuitively, two effects are at work. On the one hand, less fierce competition means that the incumbent can score relatively more wins without adopting, as the startup will quit investing already at a small disadvantage. Since paths where the incumbent wins without adoption are associated with fewer successes under the new design, having more of these paths implies that startup wins become associated with better early histories. On the other hand, less fierce competition means that games are shorter, so that the early history is more important for the final outcome. A bad early history is thus more valuable as a signal delaying adoption and startup wins are associated with worse early histories, conditional on adoption.

4.1.3. Learning versus performance racing effects

More successes in the pre-adoption phase help the startup catch up but also signal that the new design is superior and thus encourage quick adoption, which favors the incumbent. To further disentangle these two effects, Fig. 8 presents winning probabilities for the startup for the baseline case with $q_0 = 0.9$ and conditional on $T = 10$ trials. The horizontal axis measures the number of successes S and every line corresponds to a different value of the performance difference d. The vertical axis measures the startup's probability of eventually winning conditional on states (d, S, T), with $T = 10$. The figure is best read together with the top right hand panel of Fig. 5 which represents equilibrium actions in the relevant states.

A win by the startup is most likely if (i) the startup himself is just optimistic enough about the new design to continue investing, but (ii) the incumbent remains pessimistic enough to delay adoption. For example, consider the states where $S = 5$. The expected potential of the new design is only one half, and thus significantly below that of the old design, $p = 0.7$. Nevertheless, if the startup has been able to at least thrive in a niche ($d \leqslant 4$), he keeps developing the new design. As a result, his probability of winning goes from 30 percent for $d_0 = 4$ to 80 per cent for $d_0 = 0$. Moreover, for a given market share (that is, holding fixed d), the startup's probability of winning after the bumpy early history that led to $S = 5$ is higher than it would be if the expected potential of the new design were approximately equal to that of the old design ($S = 7$). It is significantly higher than it would be if the expected potential were equal to

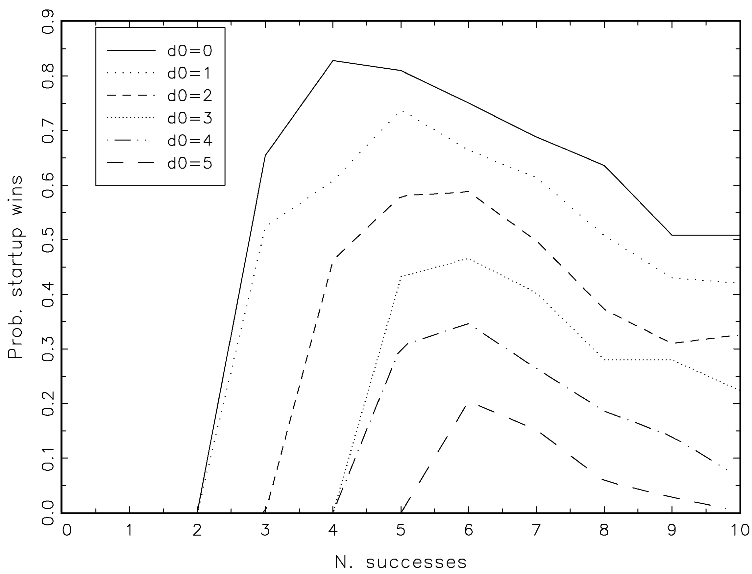

The vertical axis measures the startup's probability of winning, conditional on the state of the game (d, S, T). The number of period of play is fixed at $T = 10$. The number of successes S is measured along the horizontal axis. Different lines correspond to different values of the incumbent's performance advantage d. Parameter values are as in Fig. 6, $q_0 = 0.9$.

Fig. 8. Predicting a change in market structure.

true potential ($S = 9$). A change in market structure will thus often go along with a mediocre early history that "hides" the potential of the new design. This helps the startup to become the dominant firm by discouraging early adoption by the incumbent.

4.2. Market power, inertia and the value of uncertainty

In this subsection we show how the startup's chance of winning depends on the presence of uncertainty and the quality of the old technology. We also point out that the incumbent may lose without ever adopting. These results are established for our baseline parameterization.

4.2.1. The incumbent may lose without ever adopting

It seems puzzling that incumbents can lose market share to a startup with new design, but still never choose to adopt and eventually vanish from the market. Our model provides a rationale for such behavior, based on gambling for resurrection through a risky race with unequal designs. The third subcolumn of a broad column in Table 1 records the probability of a startup win without adoption. We emphasize that our model rules out any adjustment costs; lack of adoption comes from risk considerations alone.[13] A "bypass" outcome is more likely if the initial advantage is not too large, since the propensity to adopt is then lower. It also requires that the new design not be "too good." Again, this is because a good design will be quickly recognized by an incumbent.

4.2.2. A better old design can favor the startup

Suppose the incumbent's old design has higher known potential, holding fixed the potential of the new design. Our model predicts that, *if the new design is good enough, the startup is more likely to win against an incumbent with a better old design*. The key to this result is that an incumbent with a better old design will need to see more evidence in favor of the new before adoption. The third subcolumn of Table 1 shows that the incumbent is more likely to lose without adopting the better the old design. Table 2 shows that adoption time increases dramatically going from

[13] We have experimented with versions of the model in which adoption is costly, either in monetary terms or because it entails a loss in performance. As expected, the number of bypasses increases with adoption costs.

$p = 0.5$ to $p = 0.7$ and that the advantage at the time of adoption is lower for $p = 0.7$. This is not an obvious result, given that the longer delay for the second case is on average associated with a higher number of successes for the incumbent before adoption. The delay in adoption more than counteracts the higher average performance of the old design.

4.2.3. Uncertainty favors startups only if the new and old designs are sufficiently different

If there is no uncertainty about the potential of the new design, the incumbent always employs the best design. Given that mistakes occur under uncertainty, one might expect the startup to win more often in this case. The first subcolumn in Table 1, which reports the incumbent's winning probability for the uncertainty (first line) and certainty (second line) games, confirms this intuition only for initial conditions such that q_0 is sufficiently different from p. If the new design is inferior, preemptive adoption by the incumbent (a type I mistake) induces a symmetric race, which would never occur under certainty. If the new design is better, both failure to adopt (a type II mistake) and *delay* in adoption help the startup. In contrast, *if the two technologies are similar in quality, the incumbent is actually more likely to win under uncertainty*. A negative initial performance might discourage a startup that has fallen behind from investing, because he underestimates the actual potential of the new technology; at the same time, mistakes (both in terms of adopting an inferior technology or delaying the adoption of a superior one) are not costly for the incumbent if the technologies are similar.

4.3. Technology selection

We say that "the market selects the better design" if the winner—the firm with higher market share in the absorbing steady state—employs the design with higher potential at the end of the episode. This language is appropriate, since in the long run the majority of output is produced using the winner's design. In addition, the typical equilibrium features a leadership phase in which only the (eventual) winner invests in R&D. Thus the majority of R&D investment is also done for the winner's design.

With no uncertainty about potential, the market virtually always selects the better design. If the new design is better than the old design, this is certain: the incumbent will immediately adopt. If the old design is better, the situation is not as clear. While the incumbent never adopts, the startup is constrained to employ the new design. In the ensuing asymmetric race, the startup might win out with a worse design, especially if the difference in potential is small. However, this virtually never happens in our simulations. The reason is apparent from the left panel of Fig. 5: if the incumbent has a superior design, he will keep investing even if he has fallen very far behind. As a result, the potential of the better old design will typically carry him to victory in the long run.

4.3.1. Preemptive adoption of an inferior new design

A type I mistake (the market selects a worse new design) occurs if *either* the incumbent does not adopt the worse new design, but the startup still wins, *or* the incumbent switches to the worse design. As in the certainty case, the former is unlikely. In most cases, *selection of an inferior new design results from preemptive adoption by the incumbent*.

Table 1 shows when type I mistakes occur. Their likelihood can be read off from the columns that correspond to $q_0 < p$: it equals one minus the probability of non-adoption (subcolumn 2), plus the probability of bypass (subcolumn 3). The latter is generally negligible. Not surprisingly, type I mistakes are more likely if the difference in potential is smaller. Moreover, the likelihood of a type I mistake is hump-shaped in the incumbent's advantage. Two counteracting effects are at work here: first, given the belief, the incumbent is more likely to adopt if he is further ahead (see Section 3). On the other hand, if the incumbent is far ahead, the startup is more likely to quit investing before the belief about the new design has changed much. This lowers the probability of adoption. Overall, preemptive adoption is most likely for intermediate levels of initial advantage.

4.3.2. Suppression of a superior new design

A type II mistake occurs if the incumbent does not adopt a better new design, but still wins. The likelihood of a type II mistake can be read off from the columns in Table 1 that correspond to $q_0 > p$: it equals the probability of non-adoption (subcolumn 2), minus the probability of bypass (subcolumn 3). A type II mistake occurs if the incumbent can increase his market share enough to make the startup quit before the potential of the new design becomes apparent. Naturally, this is more likely the lower the difference in potential and the further the incumbent is ahead initially. We

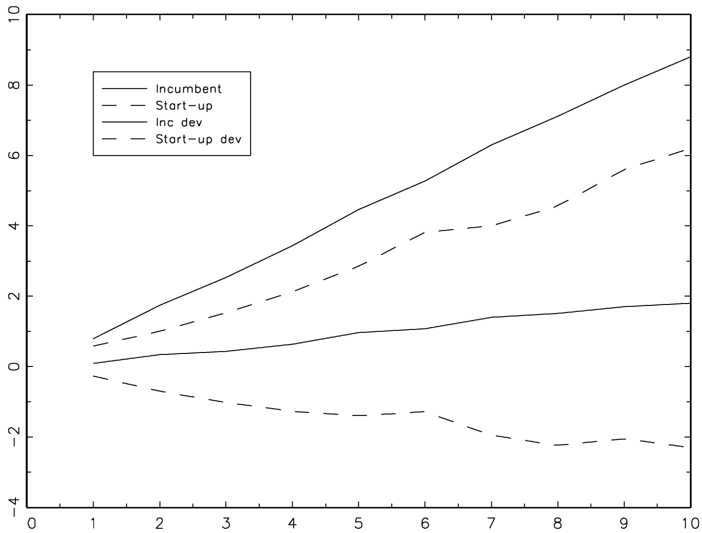

The horizontal axis represents the number of periods for which the game has been played. The top two lines describe, for each firm, the evolution of its performance level (d^I or d^S) along the average path that ends in a type II mistake. The bottom two lines describe, for each firm, the difference between its performance level along the average type II mistake path and its performance level along the (unconditional) average path. Parameter values are as in Fig. 6.

Fig. 9. The early history before a type II mistake.

conclude that, starting from a symmetric initial market structure, *a moderate advantage of the incumbent makes both types of mistakes more likely*. As the advantage increases further, mistakes occur more often if and only if the new design is better.

Figure 9 gives an idea of the transition dynamics associated with type II mistakes; it shows the average performance levels and deviations from the average long-run performance (conditional on p and q_0, respectively)[14] for both firms along histories that end in type II mistakes. Incumbents not only experience faster initial performance growth than startups, but their performance is above the long run average of the old design, whereas the opposite occur for the new design. The model thus predicts that *in episodes in which a new design is introduced, not adopted by incumbents, but ultimately driven out of the market, we should also see an unusual burst in performance of the old design right after the introduction of the new design*. We should also expect to see the old design outperform the new design.

We have experimented with different parameter configurations. The results are in line with those of the baseline case. In general, configurations associated with a longer average length of the episode increase the chances of a type I mistake (as the incumbent is more willing to preventively adopt the new technology) and reduce that of type II (as the startup is less likely to stop investing early). Both type II mistakes and the type I mistakes due to suppression of a better new technology are reminiscent of armed-bandit models of single-agent endogenous learning: actions converge to a region where no information is created before beliefs have converged to the true parameter value. However, for type I mistakes, irreversibility is a separate factor. With positive probability, we can reach a situation where a mistake was made and everybody believes this with hindsight. This situation cannot arise in an armed-bandit model.

5. Conclusion

This paper has proposed a simple framework for understanding episodes of disruptive change. We have emphasized uncertainty as a key reason for slow adoption of new technologies. Moreover, we have shown that the risk-return tradeoff that underlies an incumbent's adoption decision is very different from that in a single-firm model of technology choice under uncertainty. On the one hand, strategic competition endogenously determines firm risk attitude.

[14] For example, after T trials in the new technology, the deviation from the average long-run performance is computed as $(\frac{1}{n}\sum_{i=1}^{n} S_i - q_0)T$, where n is the number of times each episode is repeated.

On the other hand, the riskiness of an investment race depends on the conditional variance of *relative* potential. Our results rely on two key elements of the model: learning and strategic interaction.

An interesting extension of our setup would be to consider the research and development stage that precedes an episode of disruptive change. In the present paper, we have focused on the evolution of an industry *after* an initial (disruptive) innovation has been made. Nevertheless, the model has some interesting implications for the reward from innovation. In particular, a startup who introduces a better new design—in terms of expected performance growth—may expect a lower payoff. There are two counteracting effects: while a better design helps the startup gain market share more quickly, it also signals the quality of the new design to the incumbent. To exploit the first effect, but mitigate the second, it is best to enter with a design that is only marginally better than the old design. These considerations will affect why and when an entrant brings a new design to market.

Acknowledgment

We would like to thank Alfonso Gambardella, Kenneth Judd, Luigi Orsenigo, Ariel Pakes, Tom Sargent, Jean Tirole, participants at many seminars, and especially an associate editor and two anonymous referees, for helpful comments and suggestions. All remaining errors are our own.

Appendix A

In this appendix, we sketch an algorithm that has proven reliable and fast in our numerical experiments. The algorithm computes exact SMPEs, including cases that feature mixed actions in some states of the world; it does not introduce an artificial "public randomization device." The algorithm proceeds in two steps. In a first step, it finds SMPEs for games with known technologies. To this end, iteration on the operator in (5) is augmented with a "local guess-and-verify procedure" that constructs a fixed point explicitly if convergence has failed. The second step uses the fact that the Markov process in the full game has no recurrent states to reduce solving the Bellman equations to solving a sequence of low-dimensional nonlinear equations. Existence of SMPEs follows from the results in Maskin and Tirole (2001), since the game has finite action spaces and discounting of payoffs.

A.1. Step 1: Two known designs

If the potential of the new design is known, the state is simply d, together with the flag for adoption ϕ. Since ϕ can only switch from zero to one, we can solve the subgame after adoption first, and then solve a game that allows for adoption taking as given the value functions from the subgame. For simplicity, we refer to these games as "the game with $\phi = 1$" and "the game with $\phi = 0$," respectively. For both games, we know that no firm will invest if d is large enough in absolute value. We thus restrict attention to games with a finite, one-dimensional state space: we pick a large number d^{hi} and assume the state space is $\Omega^c = \{d \in Z, -d^{hi} \leqslant d \leqslant d^{hi}\}$. Accordingly, Markov strategies $\sigma^k : \Omega^c \to P(A^k(\phi))$, are restricted such that $\sigma^k(d^{hi})$ and $\sigma^k(-d^{hi})$ place probability one on no investment for all k and ϕ. For the same reason, a value function for player k is now a function $V^k : \Omega^c \to \Re$ that satisfies

$$V^k\big(d^{hi}\big) = \frac{R^k(d^{hi})}{1-\beta},$$
$$V^k\big(-d^{hi}\big) = \frac{R^k(-d^{hi})}{1-\beta}.$$

A.2. Stating equilibrium conditions using operators

We now introduce some notation and language that simplifies the presentation of the algorithm. For value functions, actions, and strategies we omit the player superscript to denote a pair of objects, that is, $\sigma = (\sigma^I, \sigma^S)$ and so on. It will be useful to order types of action pairs. For every pair of (possibly mixed) actions $a \in P(A^I) \times P(A^S)$, define an index $J(a)$ by

$$\begin{aligned} J(a) = {} & 100 I_{\{a(x^I=1)=1\}} + 10 I_{\{a(\delta=1)=1\}} + I_{\{a(x^S=1)=1\}} \\ & - 100 I_{\{a(x^I=1)\in(0,1)\}} - 10 I_{\{a(\delta=1)\in(0,1)\}} - I_{\{a(x^S=1)\in(0,1)\}}. \end{aligned}$$

The function J assigns a unique code to each pair of pure actions, as well as to each set of pairs of mixed actions with the same supports for the strategies. It will be used below to select equilibria and to formulate loops over different types of candidate equilibria.

For every $\phi = 0, 1$, we define an operator that maps pairs of value functions and pairs of Markov strategies into pairs of value functions $F_\phi : (V, \sigma) \longmapsto (W^I, W^S)$ by specifying, for $k = I, S$ and $d \neq d^{hi}, -d^{hi}$,

$$W^k(d) = \{R^k(d) + E[-cx^k + \beta \tilde{W}^k(d', \phi') \mid d, \phi, \sigma^k, \sigma^j]\} \tag{A.1}$$

where $\tilde{W}^k(d', \phi') = V^k(d')$ if $\phi' = \phi$ and $\tilde{W}^k(d', 1)$ is exogenously given if $\phi = 0$, where actions are distributed according to the strategies σ, and where the state variables evolve according to

$$\begin{aligned}
&d' = d + z^I - z^S, \\
&\phi' = \phi + \delta, \\
&Prob(z^I = 1 \mid \omega, \sigma^I, \sigma^S) = p\sigma^I(\{1, 0\}) + q\sigma^I(\{1, 1\}), \\
&Prob(z^S = 1 \mid \omega, \sigma^I, \sigma^S) = q\sigma^S(\{1\}).
\end{aligned}$$

If $\phi = 1$, then the game never reaches states with $\phi = 0$, so the specification of the value function $W^k(d', 0)$ is irrelevant. In contrast, the value function $W^k(d', 1)$ is an important component of the operator F_0.

For a given ϕ, pair of value functions V and state d, consider the one shot game with strategy spaces equal to the action spaces for that state and with payoff functions given by the right-hand side of (A.1). We can find the Nash equilibria of this game. We select unique equilibria as follows.[15] If $\phi = 0$ or $\phi = 1$ and $d \geqslant 0$, select the equilibrium with the largest index J. If $\phi = 1$ and $d < 0$, take the equilibrium for the state $-d$ and reverse the players' actions (this respects symmetry of the game with $\phi = 1$). Let $\sigma^*_\phi(V)$ denote the Markov strategy that consists of the selected Nash equilibrium actions for every state, when the continuation payoff are built from the pair of value functions V.

A pair (V, σ) is a Markov perfect equilibrium of the game given ϕ if $V = F_\phi(V, \sigma)$ and $\sigma = \sigma^*_\phi(V)$. We also define an operator G_ϕ that maps pairs of value functions into pairs of value functions by $G_\phi(V) = F_\phi(V, \sigma^*(V))$. A pair of Markov perfect equilibrium value functions for the game given ϕ is a fixed point of G_ϕ. The purpose of the operator F_ϕ is that it can be used to compute utility under arbitrary Markov strategies that are not necessarily equilibrium strategies. Indeed, we can fix a strategy $\tilde{\sigma}$ and solve for a pair of value functions W such that $W = F_\phi(W, \tilde{\sigma})$. If $\phi = 1$ we can use symmetry of the value function to reduce the number of equations and unknowns.

A.3. Action structures

The problem with standard iteration on value functions in our context is that convergence fails in states that are on the border between investment and non-investment regions. To address this, we need to perturb strategies around states where firms stop investing. To this end, it is helpful to define, given ϕ and a pair of strategies σ, the associated *action structure* as a map $\alpha^\sigma_\phi : \Omega^c \to Z$ that records, for every state, the index of the players' actions in that state under the strategies σ, that is, $\alpha^\sigma_\phi(d) = J(\sigma(d))$. The difference between strategies and an action structure is thus that, for a state where mixed actions are played, a strategy consists of probability distributions over actions, whereas the action structure only records the support of these probability distributions.

Suppose we are given an action structure $\tilde{\alpha}_\phi$. We can then use the operator F_ϕ to find strategies that respect that action structure as well a pair of value functions that represent utility under the strategy. If there are no states with mixed actions (that is non-singleton support sets), then the action structure $\tilde{\alpha}_\phi$ is equivalent to a strategy, say $\tilde{\sigma}$. We can find the value functions by solving $\tilde{W} = F_\phi(\tilde{W}, \sigma)$. For either value of ϕ, this is a linear system of $2(2d^{hi} - 1)$ equations in the $2(2d^{hi} - 1)$ unknowns $(\tilde{W}^k(d))$ that is straightforward to solve. More generally, suppose that there exist some states where $\tilde{\alpha}_\phi$ has non-singleton support sets, pointing to mixed actions. We can then expand the system of equations to simultaneously determine both the mixing probabilities in those states and the value functions: for every state in which some player chooses a mixed action, we add equations that make the player indifferent between

[15] Multiplicity of equilibria does not occur for the equilibrium solutions we have computed, but can occur during the algorithm before convergence, so a selection rule is needed there.

two actions. This adds as many equations as we need mixing probabilities, and we arrive at a (nonlinear) system of equations which still has as many equations as unknowns.

For every ϕ and strategy pair σ, we define an ordered set of *admissible candidate action structures* α associated with σ as follows. Let d_u denote the largest state d in which both firms invest for sure (assume that such a state exists; this has never been a limitation in our applications). An admissible candidate action structure follows the action structure dictated by the strategy everywhere except for the set of states $\{d_u+1, d_u+2, \dots, d_u+n\}$, where it allows other equilibria. Formally an admissible candidate action structure, indexed by integers $n>0$ and J_n, satisfies:

$$\alpha(d) = \begin{cases} J(\sigma(d)) & \text{if } |d| \leqslant d_u \text{ or } d > d_u + n, \\ J(\sigma(d_u)) & \text{if } d_u < d < d_u + n, \\ J_n \in \tilde{J}_\phi & \text{if } d = d_u + n \end{cases}$$

where the sets of possible action pairs at $d = d_u + n$ are given by $\tilde{J}_1 = \{99\}$ and $\tilde{J}_0 = \{111, 109, 101, 99, -101, -111\}$. In addition, admissibility requires that (i) if $\phi = 1$, then the action structure is symmetric, and (ii) if $\phi = 0$, then $\alpha(d) = J(\sigma(d))$ for $d < -d_u$. The admissible candidate action structures are ordered first by the number n (lower n first) and then by the index J_n (higher J_n first).

The algorithm now proceeds as follows.

- Set $\phi = 1$.
- Step 1A: Iterate on the operator G_ϕ, that is, compute $V_{t+1} = G_\phi(V_t)$, until the time τ where the difference between V_τ and $V_{\tau+1}$ (in the sense of maximal absolute value deviations across states) no longer *decreases* by more than a small number. If the difference between V_τ and $V_{\tau+1}$ at this point is sufficiently close to zero, go to Step 1E (a fixed point of G_ϕ is an SMPE).
- Step 1B: For the next-to-last iterate V_τ in Step 1A, record the Markov strategies $\sigma = \sigma^*_\phi(V_\tau)$.
- Step 1C: Loop over all admissible candidate action structures $\tilde{\alpha}$ associated with the strategies σ from Step 1B, ordered by (n, J_n).
  - Step 1C1: For every admissible candidate action structure $\tilde{\alpha}$, find strategies $\tilde{\sigma}$ that respect the action structure as well as value functions $\tilde{V}$ by solving the system $W = F_\phi(W, \tilde{\sigma})$ jointly with equations for the mixing probabilities as explained above. If there is no solution, repeat 1C1 with the next admissible candidate action structure.
  - Step1C2: Check whether $\tilde{\sigma} = \sigma^*_\phi(W)$. If yes, go to Step 1E, if not, repeat 1D1 with the next candidate action structure.
- Step 1E: If $\phi = 1$, then set $\phi = 0$. To specify the operator F_0, use the equilibrium value function of the game with $\phi = 1$ found in the previous step. Go to Step 1A. If $\phi = 0$, stop.

Errors can occur in Step 1C if the set of admissible candidate action structures is exhausted. This is a very rare occurrence in our experience.

A.4. Step 2: Backtracking the learning process

In an equilibrium of the full game, the Markov chain ω_t has no recurrent states. Indeed, the components S, T and ϕ can only grow and changes in d *alone* must be positive because they can arise only if the incumbent alone invests (using the old design); d can then not return to a lower value unless the startup invests, which involves a change in T. For a given number of trials T, let $\Omega_T = \{(d', S', T', \phi') \in \Omega\colon\ T' = T\}$. Every state $\omega = (d, S, T, \phi)$ has only two possible "successor states" in the set Ω_T, namely ω itself and $\omega' = (d+1, S, T, \phi)$.[16] For every $\omega \in \Omega_T$, the expected continuation utility on the right-hand side of the Bellman equation thus depends on $V^k(\omega)$, $V^k(\omega')$ as well as values of states in Ω_{T+1} and Ω_{T+2}.

Suppose that we know $V^k(\omega')$ and the values for states in Ω_{T+1} and Ω_{T+2}. We can now set up a system of two nonlinear equations in two unknowns, the pair of values for state ω, say (V^I_ω, V^S_ω). Consider the one-shot game with action spaces for state ω and with payoff functions given by the right-hand side of Eqs. (5), the Bellman equations of

[16] If $\phi = 0$, then an increase in ϕ is always also associated with an increase in T.

the full game, using $V^k(\omega) = V^k_\omega$. Let $H_\omega(V^I_\omega, V^S_\omega)$ denote the pair of payoffs in the Nash equilibrium of this game that achieves the largest index J. To find utilities for state ω, we solve $(V^I(\omega), V^S(\omega)) = H_\omega(V^I(\omega), V^S(\omega))$. This system is nonlinear since the equilibrium strategies depend on the continuation values.

For a large enough set of trials, uncertainty about the potential of the new design has resolved sufficiently that the posterior distribution of q given S and T will be close to being concentrated around its mean $\frac{S+1}{T+2}$. We start the algorithm at some large T^{hi}. For $T = T^{hi}$ and $T = T^{hi} + 1$, we set, for all $|d| \leqslant d^{hi}$, $S \leqslant T$ and $\phi = 0, 1$, the values $V^k(d, S, T^{hi} + 1, \phi)$ equal to the values $V^k(d)$ from the game with two known designs computed in Step 1, with the probability of success of the new design set to $\frac{S+1}{T+2}$.

The algorithm proceeds as follows.

- Set $\phi = 1$.
- Step 2A: Set $T = T^{hi}$.
- Step 2B: Let $S = T$.
- Step 2C: Set $d = d^{hi} - 1$, and let $V^k(d^{hi}, S, T, \phi) = \frac{R^k(d^{hi})}{1-\beta}$.
- Step 2D: Using the known values of states in Ω_{T+1} and Ω_{T+2} and the known value of $\omega' = (d + 1, S, T, \phi)$, form the map H_ω and find $(V^I(\omega), V^S(\omega))$ as the solution to the pair of equations $(V^I(\omega), V^S(\omega)) = H_\omega(V^I(\omega), V^S(\omega))$.
- Step 2E: If $d > -d^{hi} + 1$, let $d = d - 1$ and go to Step 2D.
- Step 2F: If $S \geqslant 0$, let $S = S - 1$ and go to Step 2C.
- Step 2G: If $\phi = 1$, set $\phi = 0$ and go to Step 2A, else stop.

Errors can go occur in Step 2D if there is no solution to the system of nonlinear equations. Again, this occurs rarely in our experience.

References

Aghion, P., Espinosa, M.P., Jullien, B., 1993. Dynamic duopoly with learning through market experimentation. Economic Theory 3, 517–539.

Anderson, A., Cabral, L., in press. Go for broke or play it safe? Dynamic competition with choice of variance. RAND Journal of Economics.

Anderson, S.P., De Palma, A., Thisse, J.F., 1992. Discrete Choice Theory of Product Differentiation. MIT Press, Cambridge, MA.

Bergemann, D., Valimaki, J., 1996. Learning and strategic pricing. Econometrica 64 (1125), 59.

Besanko, D., Doraszelski, U., 2004. Capacity dynamics and endogenous asymmetries in firm size. RAND Journal of Economics 35 (1), 23–49.

Bhattacharya, S., Chatterjee, K., Samuelson, L.W., 1986. Sequential research and the adoption of innovations. Oxford Economic Papers 38, 219–243.

Bolton, P., Harris, C., 1999. Strategic experimentation. Econometrica 67, 349–374.

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, 543–573.

Cabral, L.M., 2003. R&D competition when firms choose variance. Journal of Economics and Management Strategy 12, 139–150.

Chari, V.V., Hopenhayn, H., 1991. Vintage human capital, growth and the diffusion of new technology. Journal of Political Economy 99, 1142–1165.

Christensen, C.M., 1993. The rigid disk drive industry: A history of commercial and technological turbulence. Business History Review 67, 531–588.

Christensen, C.M., 1997. The Innovator's Dilemma. Harvard Business School Press, Boston.

Doraszelski, U., Markovich, S., in press. Advertising dynamics and competitive advantage. RAND Journal of Economics.

Doraszelski, U., Pakes, A., 2007. A framework for applied dynamic analysis in IO. In: Armstrong, M., Porter, R. (Eds.), Handbook of Industrial Organization, vol. 3. North-Holland, Amsterdam, pp. 1887–1966.

Ericson, R., Pakes, A., 1995. Markov-perfect industry dynamics: A framework for empirical work. Review of Economic Studies 62, 53–82.

Henderson, R.M., 1993. Underinvestment and incompetence as responses to radical innovation: Evidence from the photolithographic industry. RAND Journal of Economics 24 (2).

Jensen, R., 1982. Adoption and diffusion of an innovation of uncertain profitability. Journal of Economic Theory 27, 182–193.

Jovanovic, B., Nyarko, Y., 1996. Learning by doing and the choice of technology. Econometrica 64, 1299–1310.

Judd, K., 2003. Closed-loop equilibrium in a multi-stage innovation race. Economic Theory 21, 673–695.

Maskin, E., Tirole, J., 2001. Markov-perfect equilibrium: I. Observable actions. Journal of Economic Theory 100 (2), 191–219.

McCardle, K., 1985. Information acquisition and the adoption of new technology. Management Science 31, 1372–1389.

Pakes, A., McGuire, P., 1994. Computing Markov-perfect Nash equilibria: Numerical implications of a dynamic differentiated product model. RAND Journal of Economics 25, 555–589.

Parente, S.L., 1994. Technology adoption, learning-by-doing, and economic growth. Journal of Economic Theory 63, 346–369.
Tushman, M.L., Anderson, P., 1986. Technological discontinuities and organizational environments. Administrative Science Quarterly 31, 439–465.
Tushman, M.L., Anderson, P., 1990. Technological discontinuities and dominant designs: A cyclical model of technological change. Administrative Science Quarterly 35, 604–633.