DISCUSSION PAPER SERIES

No. 4925

STRATEGIC EXPERIMENTATION AND DISRUPTIVE TECHNOLOGICAL CHANGE

Fabiano Schivardi and Martin Schneider

INDUSTRIAL ORGANIZATION

Centre for Economic Policy Research

www.cepr.org

Available online at: **www.cepr.org/pubs/dps/DP4925.asp**

ISSN 0265-8003

STRATEGIC EXPERIMENTATION AND DISRUPTIVE TECHNOLOGICAL CHANGE

Fabiano Schivardi, Banca d'Italia and CEPR
Martin Schneider, New York University

Discussion Paper No. 4925
February 2005

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 4925

February 2005

ABSTRACT

Strategic Experimentation and Disruptive Technological Change*

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 firms improve both a new and a rival old technology while learning about the relative potential of both technologies. We use the model to understand historical evidence on diffusion and market structure. In particular, the model explains why a change in market leadership often goes along with slow diffusion. It also provides a rational explanation for observed 'incumbent inertia' and shows how markets can make mistakes in the selection of new technologies.

JEL Classification: C63, C73, D83, L13 and O31
Keywords: dynamic games, innovation, learning and oligopoly

Fabiano Schivardi
Research Department
Banca d'Italia
Via Nazionale, 91
00814 Roma
ITALY
Tel: (39 06) 479 22168
Fax: (39 06) 4792 820
Email: fabiano.schivardi@bancaditalia.it

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=145224

Martin Schneider
Department of Economics
New York University
269 Mercier Street, 7th Floor
New York, NY 10003
USA
Tel: (1 212) 998 8957
Email: ms1927@nyu.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=161338

*We would like to thank Alfonso Gambardella, Kenneth Judd, Luigi Orsenigo, Ariel Pakes, Tom Sargent, Jean Tirole and participants at many seminars for helpful comments and suggestions. All remaining errors are our own. The views expressed here are our own and do not necessarily reflect those of the Bank of Italy.

Submitted 02 February 2005

1 Introduction

To affect economic growth, new technologies must be both invented and widely *adopted*. Adoption of a successful new technology by an industry is typically a gradual process. At a given point in time, competing firms often employ and invest in different technologies. These facts have generated a large literature. In particular, theories of vintage capital, learning-by-doing and special interests within organizations have explained delayed adoption with firm-specific adoption costs. Diffusion of a technology is then slower when differences in adoption costs are greater. A common feature of these theories is that a new technology is immediately *known* to be superior to the old technology that it eventually – and inevitably – replaces.

In contrast, historical accounts of industry evolution emphasize that the long run potential of a new technology is often uncertain at the time the technology is first introduced. More specifically, the literature distinguishes *incremental* and *disruptive* innovations. Incremental innovations occur continually. They improve the *performance* of a product or process, and their payoffs are typically known. Disruptive innovations occur only occasionally. They change the *design* of a product or process. The potential of a new design is uncertain when it is first introduced. Moreover, the initial performance of products made under a new design is often worse than performance under the dominant old one. Introduction of a new design kicks off what has been called an *episode of disruptive technological change*. During such an episode, firms learn about the potential of the new design, as investment in incremental innovations improves both the old and the new design. Eventually, competition determines whether the new design will be widely adopted to supplant the old one.

As a concrete example, consider the hard disk drive industry, described in detail by Christensen (1997). In the 1970s and 1980s, 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 a new, smaller, drive and a dominant, larger, drive. Initially, the smaller drive had less storage capacity (that is, lower performance), and hence a smaller market share. Moreover, it was not clear whether the 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.

The historical accounts suggest that uncertainty matters for the speed of diffusion of an eventually successful new design. They also raise new questions. What does it take for a new, unproven design to actually become successful? When will a successful new design be adopted by incumbents, and when will entrants take over without incumbents adopting? Does competition always select the better of two designs? To answer these questions, the present paper proposes a strategic model of technology adoption with learning. Its key new feature is the explicit distinction between the known, but continually improving, performance of a product and the uncertain, but fixed, potential of a design used to make that product.

In our model, an episode of disruptive change is a phase where firms simultaneously learn about the potential of a new design *and* engage in a continual race to improve relative performance, thereby gaining market share. Incremental innovations not only serve as signals about the potential of the new design, but they also change firms' positions in the race. They matter for the adoption decision in both guises: new signals affects the costs and benefits of adoption, while market share affects risk attitude. A critical role is played by the *early history* of incremental innovations that occurs right after the new design has been introduced. The joint evolution of beliefs and performance during this early history determines the success or failure of a new design as well as the speed of diffusion.

In particular, the model predicts a link between diffusion and changes in market structure. Episodes that feature slow diffusion and changes in market leadership should begin with slow early performance growth under the new design. In contrast, fast early performance growth should entail fast diffusion and a stable market structure. These predictions are consistent with evidence on the cross section of episodes of disruptive change. We also derive predictions linking adoption rates, firms' subjective beliefs and market share that are supported by firm level evidence. We then show that markets can make mistakes when selecting technology: incumbents may be able to suppress good designs or preemptively adopt bad designs. We characterize how the likelihood of both mistakes depends on market structure.

Formally, our model is a dynamic investment game. Firm profits depend on (relative) performance. Two firms – incumbent and entrant – continually invest to improve performance through incremental innovation. The design employed by a firm determines how quickly performance can be improved on average: the *potential* of a design is identified with the probability that a fixed size investment project succeeds in improving performance. There are two designs: the incumbent initially employs an old design, but may choose to adopt a new design that is being employed by the entrant.

It is initially unknown whether the new design has more potential than the old one. Firms learn about this in a Bayesian fashion by observing each other's progress. The incumbent's adoption decision thus becomes a choice between two "races", one with equal designs (that can be entered through adoption) and one with unequal designs (run as long as there is no adoption). Like any investment decision under uncertainty, it depends on the relative risk and return of the choices as well as firms' risk attitude.

A race with equal designs is less risky, since there is no difference in potential between the two firms, even if the potential of the design itself is unknown. It follows that adopting the *unknown* new technology actually *reduces* risk for incumbents. This result stands in sharp contrast to what one might conclude from a a single-firm problem. It emphasizes that any assessment of risk must take the strategic context into account.

Firms' risk attitude depends on their position in the race. When market size is finite, market leaders have diminishing marginal gains from better performance, and act in a risk averse manner. In contrast, laggards with small market share face increasing marginal gains and hence are risk-loving.[1] It follows that incumbents with higher market share are more likely to adopt: they prefer to reduce risk by moving to the race with equal designs. In contrast, 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. As in the data, 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

[1]Formally, firms are risk neutral players, as is standard in oligopoly models. 'Risk attitude' refers to the curvature of the value function and is determined endogenously. We show that firms' value functions in a Markov Perfect Equilibrium are convex when relative performance is low and concave when it is high

the entrant 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 fits well with evidence on the cross-section of episodes, discussed in Section 5 below, where we review the history of the earth-moving equipment and hard disk drive industries. In particular, episodes characterized by "incumbent inertia" – a dramatic loss in incumbent market share accompanied by a failure to adopt a new technology – are often preceded by a phase during which the new technology initially languishes in a market niche for some time. While the management science literature explains incumbent inertia with naive extrapolation on the part of incumbent management, our model points to strategic decision making under uncertainty: we show that it can be optimal not to adopt (costlessly) even if the expected potential of the new design is higher.

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, the incumbent may *preemptively adopt* an inferior design. If the new design is expected to do as well as the old design, but there is still considerable uncertainty about relative potential, it makes sense for a market leader to adopt simply to reduce the risk of racing with unequal designs. Of course, learning will generally continue after adoption. As a result, we may end up in a situation where everybody agrees that the new design turned out to be inferior. On the other hand, the incumbent may *suppress* a superior new design. Indeed, if the latter design performs relatively badly during the early history, the incumbent may reestablish dominance without adopting. Since learning ceases once the incumbent has pulled ahead far enough that the entrant stops investing in the new technology, neither firm will ever realize that the better design is actually languishing in a niche. 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 entrant, 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 are a fixed point of a generally *nonmonotonic* operator on tuples of value functions, making value function iteration inapplicable. In addition, we have to deal with both beliefs and performance as state variables.[2] We solve the first problem by using a combination of value function iteration and a "local guess-and-verify" procedure which 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 converged, and then backtracking from there, thus avoiding costly iteration. Throughout, we solve for exact equilibria, without the need for a public randomization device.

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 discusses the results in light of existing empirical evidence. Section 6 concludes.

Related Literature

The paper is related to several strands of literature. The large body of work on patent races is also concerned with whether incumbents or entrants are more innovative. However, patent race models focus on the R&D stage that takes place before a product is brought to market. They also assume that the prize to the race winner is known. In contrast, we are interested in competitive situations where entry has already occurred, but where the potential of a new technology is not yet known. Patent race models cannot be used to study such situations, because they do not allow for incremental innovation and learning.

If learning and technology choice were absent, our setup would reduce to a multi-stage investment game of the type analyzed by, for example, Fudenberg et al.(1983); Judd (1985), Harris (1989), Dutta and Rustichini (1996), Budd, Harris and Vickers (1993), or

[2]Existing algorithms for investment games (Ericson and Pakes 1995, Pakes and McGuire 1994) do not allow for learning and technology choice.

Aghion *et al.* (2001). In fact, the subgame that firms in our model enter *after* adoption is essentially a discrete-time version of the game studied by Budd, Harris and Vickers. A key principle identified by these authors is that competition tends to move firms towards states in which joint payoffs are higher. In our model, where total profits increase with the market leader's performance advantage, this principle gives rise to an "increasing dominance" result, i.e. the leader tends to move further and further ahead. We use this to model "episodes" that end in finite time.

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 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 (2004) use a similar framework to study the size distribution of firms and the role of advertising, respectively. As the papers on multi-stage patent races, they are interested in establishing conditions under which increasing dominance occurs. Our approach differs from all these papers in that we consider technology choice and learning.

Delayed adoption has 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 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 creates uncertainty about how to operate the new 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.[3] 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.

Our setup implicitly incorporates firm-specific adoption costs for incremental innovations, since the effect of investment on performance is stochastic. Since investment projects

[3]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.

may fail to improve performance, firms can differ in performance even if they employ the same design. In other words, it is not costless for a firm to implement an incremental innovation even if another firm has already achieved it. This captures firm-specific capabilities to undertake innovations, as in Jovanovic and MacDonald (1994). A large literature in management science provides evidence that organizational capabilities matter for adoption decisions (for example, Tushman and Anderson 1986, 1990, or Henderson and Clark, 1990). Tornell (1997) presents a model of special interests within organizations that try to block adoption.

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, Chatterjee and Samuelson (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 in our model 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.[4]

[4]Existing literature on strategic experimentation (Aghion, Espinosa and Jullien 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.

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 and there is an infinite horizon. 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 Appendix A.1, 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. Figure 1 plots two examples of symmetric revenue functions obtained by solving the Bertrand pricing game.

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 employs. There are two technologies, 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 old firm initially employs the old design, but may irreversibly switch to the new design at any time. To single out the role of learning, we assume that adoption entails no direct costs for the incumbent. We denote the adoption choice by $s_t \in \{0, 1\}$ and let

$\phi_t = 1$ indicate whether adoption has occurred prior to time t. Firms' action spaces are

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

$$A_t^I(\phi_t) := \{(x_t^I, s_t) \in \{0,1\}^2 : x_t^I \geq s_t \text{ and } s_t \leq 1 - \phi_t\} \tag{2}$$

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.

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

$$h^t = (x_\tau^I, x_\tau^S, z_\tau^I, z_\tau^S, s_\tau; \tau = 1, 2, ..., t-1)$$

Let H^t denote the set of period t histories. A *strategy* for player k is a collection of functions $\sigma^k = \{\sigma_t^k; \sigma_t^k : H_t \rightarrow P(A_t^k(\phi^t)); t = 0, 1...\}$; each σ_t^k, 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

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

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

where relative performance evolves, for given d_0, according to

$$d_{t+1} = d_t + z_t^I - z_t^S; \ t \geq 1 \tag{4}$$

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

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, 1994). 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* $\Omega = \{(d, S, T, \phi) \in Z \times Z_+^2 \times \{0, 1\} : S \leq T\}$.

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

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

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

s.t.

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

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

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

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

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

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

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) 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, we have

Proposition

For any SMPE, there exist $d^{hi}$and d^{lo} such that for all $d \geq d^{hi}$ or $d \leq d^{lo}$ and for all $S,T \in Z_+$, $S \leq T$, and $\phi \in \{0,1\}$, $\sigma^{I*}(d,S,T,\phi)$ assigns probability one to the actions $x^I = 0$ and $s = 0$ and $\sigma^{S*}(d,S,T,\phi)$ assigns probability one to the action $x^S = 0$.

The set of states ω with $d \geq d^{hi}$ or $d \leq d^{lo}$ consists of *steady 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 steady states is *absorbing*, and 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 Computation

To our knowledge, there does not exist an algorithm for the computation of Markov Perfect Equilibria that is guaranteed to converge. In this subsection, we sketch an algorithm that has proven reliable and fast in our numerical experiments. A detailed description is provided in Appendix A.2. 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".

A problem with computing SMPE is that the operator between pairs of Bellman equations implicit in (5) is typically not monotonic. Straightforward iteration of this operator may not lead to a fixed point. Our algorithm is based on three observations. First, iteration does produce a pair of functions that does not change under the operator *except in certain bounded regions of the state space*. Second, the continuation game for high T will be very similar to a game with two known technologies, as the posterior distribution becomes concentrated around its mean. Third, the presence of learning implies that the state process has no recurrent states.

The algorithm proceeds in two steps. In a first step, the algorithm finds SMPEs for games with known technologies, i.e. with probabilities of success known in advance. 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 two features of the learning process. First, for high T, the posterior distribution of q given S and T will be close to being concentrated around its mean $\frac{S+1}{T+2}$. We thus take the value functions of the "certainty games" computed in step 1 as an approximation for the value function at high values of T.

Second, the Markov chain ω_t has no recurrent states: 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 cannot return to a lower value unless the startup invests, which involves a change in T. This allows us to backtrack the learning process by solving a sequence of low-dimensional nonlinear equation systems.

2.5 Parametrization

The revenue function used in the analysis below is the solid line in Figure 1. In Appendix A, this revenue function is derived from a model of Bertrand competition. We let performance correspond to product quality, for which consumers have logit demand. Since every consumer buys exactly one unit of the good, performance only matters for the consumers' choice of brand. As a result, firm revenues only depend on the difference in performance.

In applications, it is helpful to think of d_t^k as an index that combines several attributes of the product in question. For some industries, empirical studies have constructed such indexes (see Foster (1986) and the references therein). We return to this interpretation below when discussing specific industries. However, our Bertrand setup offers a second line of interpretation: the performance difference d maps one-to-one into market share. The dynamics of d may thus be interpreted as the evolution of market structure. In particular, the set of steady states (that is, states with $d \geq d^{hi}$ or $d \leq d^{lo}$) can be viewed as a set of "robust" market structures where a low performance and a high performance provider have split the market and have no incentive to change this situation.

We present numerical results based on a specific parametrization of the model, with parameter values listed in Figure 2. This parametrization is chosen to illustrate key qualitative properties of the model. It is representative of many parametrizations that we have computed. In what follows, we also point out the role of parameter values and functional form where we have found them to be important. In particular, we emphasize the role of learning by comparing our model to a benchmark where the potential of the new design is known.

In general terms, the qualitative results are similar for all within-period revenue functions such that the *total* profits of the two firms are bounded and increasing in the difference in performance between leader and follower. Boundedness implies that the revenue function is concave for the leader and convex for the follower; this property is then inherited by the value function and generates the risk tradeoffs we describe below. It also ensures that, if the difference in performance is large enough, it is no longer optimal for firms to invest, that is, there exists a steady state.

When total profits are also increasing, the steady state is *absorbing*, which makes our model one of an episode: a market structure with one dominant firm is reached after a finite (but random) number of periods. This reflects an important unifying principle in models of dynamic competition, identified by Budd, Harris and Vickers (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 in light of the episode evidence we are trying to explain. However, our setup and computational techniques could also be used to study ongoing competition in future work.

3 Characterizing Equilibria

This section describes equilibrium strategies and payoffs in two steps. Subsection 4.1 considers the subgame that begins when the incumbent adopts the new design. This game is a *one-technology race*, during which only incremental innovations are undertaken. Subsection 4.2. then characterizes the pre-adoption phase, where the two designs compete head-on.

3.1 The One Technology Race

In the one-technology race ($\phi = 1$), firms decide only whether to invest or not. They 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.

How Episodes End

Since $R^I(d) = R^S(-d)$, the 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. Figure 2 shows how equilibrium strategies depend on d, together with the value function. There are three equilibrium regions. By Proposition 1, for d high enough both firms will not invest. If the cost of investment is not too high, both firms invest at zero. For 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 Figure 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.

R&D Productivity and Market Structure

The parameter q_0 measures the average number of successful incremental innovations per investment project – a measure of R&D productivity. Figure 3 illustrates how R&D productivity shapes market structure in our setting. In the bottom panel, the performance advantage of one player, say the incumbent, is measured along the x-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, and the bottom left half that of the startup. In both cases, a dark/blue triangle indicates investment, while a light/gray 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 role played by the parameter q_0. On the one hand it governs the *growth* of performance in absolute terms: an episode in which performance *levels* grow faster on average corresponds to a higher value of q_0. On the other hand, q_0 determines the random walk behavior of d_t: 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 second effect: during the "Schumpeterian" phase, intermediate values for R&D productivity (i.e. q_0 close to 1/2) imply 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.

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 panel 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 x-axis now measures the performance advantage of the race leader.

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 box 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. In addition, every row of rectangles within a box corresponds to a value of S, the number of successful investment projects under the new design. Lines higher up within a box 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 the benchmark game where q is known. For every row corresponding to a pair (S, T) we have set the probability q_0 equal to the posterior mean, $\frac{S+1}{T+2}$. The basic pattern observed in Figure 3 is also present in the uncertainty case: the investment regions narrows 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 two effects – growth and variance – 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 Design Competition

We now examine the incumbent firm's adoption decision. Figure 5 shows equilibria in the pre-adoption states ($\phi = 0$), for the games with known q (left panel) and unknown q (right panel), respectively. The basic structure of these plots is the same as in Figure 4. However, 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 for the old design, and gray/green for the new design. White triangles indicate no investment.

Immediate Technology Choice under Certainty

As in the one-technology race, there is a region around $d = 0$ where both firms invest, and a no-action region when $|d|$ is high enough. The new element is the adoption decision. When q is known, this decision is easy enough: the incumbent *immediately* adopts the new design if q_0 is higher than p; he never adopts otherwise. We conclude that a model without uncertainty cannot explain *delayed* adoption by the incumbent.

For $q_0 \geq p$, the equilibrium dynamics are thus the same as in the one-technology race detailed above. For $q_0 < p$, the dynamics has a different flavor: the process d_t now has a positive drift as long as both firms invest. In theory, this does not mean that the startup cannot win the race with a worse design. However, this will not turn out to be very likely in the simulations below. To properly capture design competition, we now turn to the case where the potential of the new design is uncertain.

The Propensity to Adopt while Learning

The equilibrium pattern in the uncertainty game can be summarized by a simple rule-of-thumb: *The incumbent adopts the new design if the posterior mean is high enough or if his performance advantage is high enough.* The first property is simply that the incumbent will adopt the new design once he has seen enough evidence in its favor. Indeed, the threshold value of the posterior mean $\frac{S+1}{T+2}$ that induces the incumbent to switch decreases with the sample size T, reflecting the greater confidence the incumbent can have in his estimate.

The second property is less obvious: why should an incumbent who is further ahead in terms of performance have more of an advantage to adopt the new design? The answer follows again from the incumbent's behavior toward 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 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 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.

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, such as initial market structure (d_0) and the relative potential of the technologies (q_0 and p). They also deliver predictions about the transition dynamics. The main statistics are reported in Table 1, based on 500 repetitions of each episode.

4.1 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 Figure 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.

Preemptive Adoption of an Inferior New Design

With uncertain potential, the market can make two types of "mistakes": it can select a worse new design (a "type I mistake"), or it can fail to select a better new design (a "type II mistake"). A type I mistake 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 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.

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 can 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 6 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)[5] for both firms along histories that end in type II mistakes. Incumbents not only experience faster initial performance growth than startups,

[5]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.

but their performance is above the long run average of the old design, whereas the new design has a subpar performance. 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.

The possibility of type II mistake result (and also the second variety of type I mistakes) is 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.

4.2 The Persistence of Market Power

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.

Does Uncertainty Favor Startups ?

The first subcolumn in Table 1 confirms the above intuition, but 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 technologies or delaying the adoption of a superior one) are not costly for the incumbent if the technologies are similar.

Design Competition: Delay and Failure to Adopt

Delayed adoption of a superior design can be costly for incumbents even when they do not make a type II mistake. This is illustrated in Table 2. For example, with

$p = .7, q_0 = .85$, and an initial advantage of $d_0 = 3$, an incumbent adopts 94% of the time, and never loses after not adopting. Still, his winning percentage is only 72%. The reason is that, on average, adoption takes place after 7.15 periods. During lengthy design competition, a startup with a better design is bound to catch up: at the time of adoption the advantage has been reduced to just $d_t = 1.94$, on average.

The third subcolumn in Table 1 shows that, under uncertainty, it is possible for the incumbent to lose without having adopted the new design. *Even without adjustment costs,* the model thus delivers a rational explanation of an incumbent being "bypassed" by a startup.[6] This outcome is more likely if the initial advantage is not too large, since the propensity to adopt is then lower. A "bypass" also requires that the new design not be "too good". While a good design makes it easier for the startup to catch up during design competition, its potential will also be quickly recognized by an incumbent, cutting design competition short.

Predicting Changes in Market Power

We now consider the relationship between the *early history* of incremental innovations and changes in market structure. We first view the model from the perspective of an observer who compares whole episodes - here full equilibrium sample paths produced by the model. Suppose the observer compares the pre-adoption history for episodes in which the startup wins versus episodes in which the incumbent wins. Figure 7 shows that the early performance of the new design is better on average in the episodes in which the incumbent wins. More specifically, the dashed line in the figure is the average performance level of the startup on paths along which the incumbent has not yet adopted, but eventually wins, stated as a deviation from the unconditional average performance path. The solid line shows the same deviation for paths on which the startup eventually wins. Paths that lead to incumbent wins thus feature unusually good early performance by the startup, and in particular, better performance than paths that lead to startup wins. [7]

This finding reflects again two counteracting effects: on the one hand, more suc-

[6] 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.

[7] Both lines exhibit a large drop in period 4. This is because paths along which the incumbent has not adopted by period 4 must have experienced at least one failed investment project under the new design in the preceding three periods.

cesses under the new design help the startup catch up. On the other hand, more successes signal that the new design is superior and thus encourage quick adoption, which favors the incumbent. To further disentangle these two effects, Figure 8 presents winning probabilities for the startup after 10 trials under the new design. Here 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 winning conditional on states (d, S, T), with $T = 10$. The figure is best read together with the top right hand panel of Figure 5 which represents equilibrium actions in the relevant states.

A win by the startup is most likely if (i) the startup itself 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, $p = .7$. Nevertheless, if the startup has been able to at least thrive in a niche ($d \leq 3$), he keeps developing the new design. As a result, his probability of winning is at least 40 percent.

In particular, 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 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.

4.3 The "Incumbent Inertia" Hypothesis

The management science literature has argued that "more established firms" are more likely to be replaced if a better new design comes along (Tushman and Anderson, 1986). We now reexamine this issue in our framework. We consider whether firms with higher initial market share or a better old design are more likely to lose. Without uncertainty about the potential of the new design, it is clear from Table 1 that both types of advantage make it more likely that the incumbent will remain the dominant firm at the end of the episode. If "inertia" is to reverse this result, it must be that more established firms cannot deal as well with uncertainty.

High Initial Market Share Implies Incumbent Activism

Suppose the incumbent has a higher market share at the time the startup enters. Section 3 has shown that the propensity to adopt increases with market share. There is thus no force that creates inertia. Not surprisingly, Table 1 shows that *the incumbent's winning percentage is increasing in* d_0. In addition, the difference in winning percentages between the certainty and uncertainty games decreases as d_0 increases: *uncertainty is more harmful for incumbents with low initial advantage in performance.*

The greater success of leaders can be traced to adoption behavior. The adoption percentage is hump-shaped in d_0. For low values, the incumbent is likely to fall behind and "gamble for resurrection"; for high values, he is likely to win early on, before the true potential has been revealed. Table 2 shows that the average delay in adoption decreases with d_0, and that the average advantage at adoption increases: an incumbent with a high advantage needs a smaller number of successes to "be convinced" of the opportunity to adopt.

High Initial Potential Implies Incumbent Inertia

Suppose the incumbent's old design has a higher probability of success, holding fixed the expected potential of the new design. In this experiment, a more established firm is thus taken to be one which is perceived, ex ante, to have a better design. Table 1 shows that, *conditional on the true potential of the new design, the startup is more likely to win against an incumbent with a better old design.* This provides a precise statement of a version of the incumbent inertia hypothesis.

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 it the incumbent is more likely to lose without adopting the better the old design. Table 1 shows that the percentages of adoption are uniformly lower and Table 2 shows that adoption time increases dramatically going from $p = .5$ to $p = .7$. A comparison between the advantage at the time of adoption shows that this is lower for $p = .7$. 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.

5 Empirical evidence

In this section, we review empirical work that relates to the main predictions of our model. We begin by discussing quantitative evidence on the link between subjective beliefs, market share and the propensity to adopt. We then review industry histories that illustrate how the early history of incremental innovations shapes the evolution of technology and market structure during an episode of disruptive change. Finally, we discuss evidence on technology selection by markets.

Beliefs, Market Share and the Propensity to Adopt

Our model predicts that an incumbent firm's propensity to adopt is increasing in its *subjective* assessment of the potential of a new technology. To evaluate the relevance of this effect, it is necessary to construct empirical measures of subjective beliefs.[8] Kaplan, Murray and Henderson (2003) consider the role of managerial beliefs in the pharmaceutical industry during the biotech revolution. Biotechnology fits our description of disruptive innovations with uncertain potential: "...the emergence of biotechnology was not a single, well-understood event, but rather a complex mix of scientific, technical and business mode changes that unfolded over several distinct phases" (Kaplan et al., 2003, pp.206-207). They analyze a sample of 15 US and UK pharmaceutical firms for the period 1973-1998. As measures of adoption, they use indicators of biotech-related output (patents and publications); for beliefs ("recognition" in their terminology), they use counts of biotech-related words in the annual letter to shareholders. They argue that this measure is a good proxy of managements' assessment of biotechnology: managers that have a higher estimate of the potential of this new technology are more likely to mention it in the letter to shareholders. They find a very robust correlation between word counts and incumbents' adoption propensity, even after controlling for a number of firm level features. We interpret this as strong evidence of the relevance of subjective beliefs.

Our model also predicts that an incumbent firm's propensity to adopt is increasing in its performance advantage. In the oligopoly structure we have used above, relative

[8]This is especially important since standard empirical analysis used for rational expectations models cannot be applied to learning models such as ours. For example, our model says that, in episodes of learning, expectations need not appear "correct" with hindsight. Instead, they start out at an – unobservable – prior and evolve over time as managers learn information not necessarily observed by researchers.

performance translates directly into market share. It is thus not necessary to define and measure performance directly – the model predicts a positive relationship between adoption rate and market share. Kaplan et al. (2003) confirm the existence of such a relationship, holding fixed their measure of managerial beliefs. This result is in line with previous findings of Zucker and Darby (1996) on the biotech revolution. Karshenas and Stoneman (1993) also find that adoption of computer numerically controlled machine tools in the UK engineering industry at the beginning of the 80s' was positively influenced by size, measured in terms of number of employees. This is also consistent with the model, provided that size is taken as a proxy for market share. Our model suggests that the higher adoption propensity is due to the higher degree of risk aversion that firms with a high market share exhibit in equilibrium.

The Early History and the Evolution of Market Structure

The management science literature that has surveyed a number of industries during episodes of disruptive change.[9] Case studies are particularly informative for our purposes, since they often contain a description of both the early history of technological competition between designs and the evolution of beliefs. The main message emerging from this body of work is that uncertainty plays a central role during episodes of disruptive change. Here we consider two industries in detail – earth-moving equipment and hard disk drives. Both have experienced several episodes of disruptive change, with very different patterns for the early history of incremental innovations.

Case 1: Disruptive Change in the Earth Moving Equipment Industry

The history of the earth moving equipment industry illustrates how two disruptive innovations can lead to very different adoption patterns and changes in market structure, as the information content of the early history varies. To map technological progress in excavators into the setup of the model, we take "performance" to be a mix of bucket capacity, maneuverability and reliability, while "design" corresponds to the technology to power and maneuver the bucket. We describe two new designs that triggered episodes of disruptive change: the gasoline engine and hydraulic-actuated machines. First, consider

[9] See for example, Tushman and Anderson (1986) for the cement, glass and minicomputer industries, Henderson and Clark (1990) for the semiconductor photo-lithographic alignment equipment industry and Iansiti and Khanna (1995) for the mainframe computer industry.

the gasoline engine, introduced in the early 1920s when existing earth-moving equipment was powered by steam. Very quickly, gasoline-powered equipment became more powerful and reliable and before the end of the decade the steam engine was replaced. Even though technology changed dramatically during this episode, market structure changed little – 23 of the largest 25 makers of steam shovels were able to successfully move to the production of gasoline-powered equipment (Christensen 1997).

A second disruptive innovation hit the excavator industry after World War II: new hydraulically-actuated equipment was introduced as an alternative to then dominant cable-actuated machines. The resulting early history of incremental innovations was quite different from the experience of the 1920s. Indeed, early hydraulic shovels were too weak for major earth moving tasks and were used only by residential contractors to dig small trenches for water and sewer lines. They were welcome for the latter tasks because they were more mobile than cable-actuated machines. Figure 9, reproduced from Christensen (1997), shows the evolution of the capacity of the new design with respect to the contractors' requirements. In the early phase of development of the new design, capacity was limited and its development was slow, so that it was not used by general excavation contractors, that represented the bulk of the market and that used cable-activated equipment. Most incumbents, therefore, kept producing and developing cable-activated equipment. For 10 to 15 years after its introduction, the development of capacity was slow, and hydraulic excavators were produced for the niche market of residential contractors. After this phase, however, the capacity of the new machines was increased at a faster pace than that of cable excavators, so that they became competitive also for heavy applications, such as mining contractors. At this point, additional advantages of hydraulic machines, particularly maneuverability and reliability, shifted the balance of performance in favor of the new design.

The slow initial development of hydraulic excavators was accompanied by a dramatic change in market structure. At the end of the episode – when hydraulic excavators had become the dominant design – only one of the 30 or so established manufacturers of cable-actuated earth moving equipment had made the transition to the production of hydraulic excavators. The bulk of the market was taken over by Caterpillar – an entrant early on in the episode – as well as a number of farm equipment makers. While a few incumbents continued to produce exclusively very large cable actuated excavators for mining, most incumbents were simply driven out of the market.

Our model explains the stark difference between the two episodes of disruptive change by the difference in the corresponding learning processes. With hindsight, the new design turned out to be superior in both episodes. However, the potential of the gasoline engine was revealed more quickly by the early history than was that of hydraulic shovels. As a result, the incumbents of the 1920s quickly adopted and market structure changed little.[10] In contrast, the incumbents of the post war periods were uncertain about the potential of hydraulic shovels and thus did not adopt quickly. Caterpillar built a small market share initially by selling its shovels in a niche. Later on, as the performance of hydraulic machines increased, the incumbents were still reluctant to switch designs. They gambled for resurrection by maintaining a risky race of unequal designs, but eventually lost on that gamble.

Case 2: Disruptive Change in the Hard Disk Industry

In the hard disk drive industry, the innovation process can also be represented along the two dimensions of our model. "Design" now corresponds to the diameter of the disk, that was progressively reduced from 14" to 3.5". "Performance" corresponds to a mix of storage capacity – which has been increasing steadily over time – and size, which plays a minor role. Over the last thirty years, market structure has been remarkably stable whenever the industry underwent phases of improvement of the storage capacity for a given design. However, every new generation of disk was introduced by entrants that eventually became market leaders.

Using data from the industry publication *Disk/Trend Report* as well as interviews with over sixty executives in the industry, Christensen (1993) offers a detailed study of the introduction of the 3.5" drive. This drive was first brought to market by a start-up, Conner Peripherals, at a time when the 5.25" drive was the dominant design. Initial performance of the new, smaller design was low, especially because of its low storage capacity. Conner was able to build a small market share by selling its drive to laptop makers. For the desktop

[10] A feature not directly captured by the model is that the gasoline engine was introduced in many industries simultaneously in the 1920s. As a result, managers probably obtained additional signals from outside their own industry. A learning view of adoption would suggest that such informational spillovers would have further increased the speed of diffusion.

market, the development possibilities were estimated to be below those of the 5.25" design. The market leader in 5.25" drives, Seagate, thus decided not to market a 3.5" model and instead continued to invest in improving the 5.25" drive.

However, over time relative performance growth of 3.5" drives was surprisingly high. On the one hand, storage capabilities increased more quickly than expected. On the other hand, there were other advantages in the smaller design that had not been initially foreseen. In particular, "less mass in the 3.5 inch drive meant less vibration and less inertia, enabling drive makers to position the head more accurately [...]. These engendering implications had not been apparent when the 3.5 inch architecture was first presented to manufacturers." (Christensen, 1993, p. 567). The 3.5" drive became able to successfully compete for the large market for desktop computers. When Seagate finally decided to switch to the new design, Conner had already accumulated enough advantage to become a major player in the industry. As for the above case of hydraulic excavators, our model explains this episode in terms of slow initial development of the new design, that delayed the incumbent's adoption decision.

The example of the 3.5" disk drive is interesting not only because it provides another instance of "incumbent inertia", but also because there is additional evidence that speaks against traditional explanations based on firm-specific adoption costs or capabilities. Indeed, the 3.5" drive had originally been developed in Seagate's research department and the start-up Conner was founded by a group of Seagate engineers.[11] This suggests that Seagate's delayed adoption of the 3.5" design was a business decision based on management's subjective expectations, not a decision forced on it by lack of capabilities.

Selection

As a recent example of Type I mistakes – a worse new technology is preemptively adopted – consider merger wave between entertainment and telecom companies in the late 1990s and early 2000s, including the Viacom/CBS and AOL/Time Warner mergers. These mergers were driven by a high estimate of the impact of fast data transmission on the distribution of entertainment to households, replacing old designs such as VHS and DVD. The strong performance of Internet companies in the second half of the nineties generated

[11] See Franco and Filson (2002) for a model of technology diffusion through employee mobility applied to the disk drive industry.

an overestimate of the potential of this new technology, prompting “old economy” firms to adopt it. In particular, in the case of the AOL/Time Warner merger, the deal was motivated by the possibility to sell Time Warner products to AOL customers.

Time Warner was one of the major players in the entertainment industry. According to our model, this made it more eager to adopt the new technology early on in order to minimize the risk of missing a disruptive development in the industry. Ex post, it is becoming apparent that the synergies between the two activities are lower than initially expected, and that many of these mergers were unjustified. In particular, the AOL/Time Warner merger has caused one of the largest destruction of value in corporate history. Rather than simple mistakes on the CEOs side, our model suggests that these events can be explained by a particularly strong initial performance of the new technology, that led to an overestimate of its actual potential, coupled with the high propensity to adopt of the incumbent firm induced by market leadership.

6 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. 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.

We have also shown that a model of rational firms operating in an uncertain environment can rationalize seemingly puzzling “inert” behavior exhibited by incumbents in many episodes of disruptive change. In particular, incumbents become more risk-loving as their market share is eroded. But more risk-loving firms are less willing to adopt a new design: they prefer the riskier race with different designs. Incumbents may thus lose market leadership without having adopted a new technology that is widely regarded to have higher potential. Such behavior is more likely when the new design got off to a slow start – a subpar early history – which is consistent with historical accounts.

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, an entrant 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 entrant 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. In particular, the payoffs from innovations that will be improved in ongoing competition after introduction are not symmetric for incumbents and entrants, contrary to what is often assumed in the literature on patent races. In addition, the optimal timing of introduction is a nontrivial problem that deserves further study.

A Appendix

A.1 Pricing and Revenue Functions

In this subsection we specify the revenue functions R^i used in our numerical work from within-period Bertrand competition between firms, given relative performance d. We adapt the standard logit model of demand (e.g. Anderson et al. (1992)). In every period there is a continuum of (one period lived) consumers of mass 1 who would like to purchase one unit of an indivisible good, as long as it is available at a price below some reservation price, say y. The incumbent and the startup firm offer different brands of this good. Consumer preferences are represented by (conditional) indirect utility: the utility obtained by consumer m from purchasing a unit of the brand produced by firm k is

$$U_m^k = \begin{cases} y - p^k + d^k + \epsilon_m^k & \text{if } p^k \leq y \\ -\infty & \text{otherwise} \end{cases}$$

where p^k denotes the price set by firm k, and the ϵ_m is independent across consumers and doubly exponentially distributed. Consumers also have an "outside option":[12]

$$U_m^0 = \begin{cases} 0 & \text{if } p^I, p^S \; > \; y \\ -\infty & \text{otherwise} \end{cases}$$

They choose a brand according to

$$U_m = max_{k \in \{0,I,S\}} U_m^k$$

Integrating over consumers we obtain market demand for firm k given the price of firm $j \neq k$:

$$X^k = \begin{cases} 1 & \text{if } p^j > y \geq p^k \\ \frac{exp(\frac{-\, p^k}{\mu})}{exp(\frac{-\, p^k}{\mu}) \; + \; exp(\frac{d^j \, - \, d^k \, - \, p^j}{\mu})} & \text{if } p^I, p^S \leq y \\ 0 & \text{if } p^k > y \end{cases}$$

Demand depends only on the difference in performance. The parameter μ governs the substitutability of products.

Firms face equal constant marginal costs $\bar{c}$ of producing output. The first order condition for the price of firm k is

$$p_k \quad \geq \quad \bar{c} + \frac{\mu}{1 \; - \; \tilde{X}_k} \quad ; \text{with equality if} \, p_k \; < \; y \tag{10}$$

$$p_k \quad \leq \quad y \tag{11}$$

These condition characterize Bertrand equilibrium prices in state d. By symmetry, the corresponding within-period profits can be summarized by a single function, say $\Pi(d)$. It is plotted in Figure 1 for different values of μ. The revenue functions to be used in the dynamic investment game are $R^I(d) = \Pi(d)$ and $R^S(d) = \Pi(-d)$, respectively.

We have defined SMPEs only in terms of the investment game. However, the fact that we solve separately the pricing game and the investment game is not restrictive. Suppose we were to consider an extended dynamic game in which firms choose prices at each point in time *simultaneously* with investment and technology. In any SMPE of the extended game, firms would still play the Bertrand pricing strategies in every state. This follows directly from time separability of intertemporal profits. Consider the Bellman equations (5): If R^k were replaced by the Bertrand profit function, the payoff function of the one shot game implicit in the right hand side of the Bellman equations would be the sum of two terms, with the price strategies determining the first (current profit) part and the investment (and switching) strategies determining the second (future expected profit) part. But then the one shot game may be viewed as constructed from two games, played simultaneously and with the payoffs added up to yield the stage game payoffs. Clearly, then, a Nash equilibrium

[12]This differs from the more conventional concept of an outside option which yields some fixed utility independent of the prices of the indivisible goods (see Anderson et al. (1992)). Demand in this case depends on the levels of performance relative to the utility provided by the outside option, and not just on the difference in performance levels. In our model, consumers want to purchase the indivisible good regardless of how high the level of performance is, provided they can afford it.

of the one shot game must be a Nash equilibrium of every single component game.

A.2 Computational Algorithm

Step 1: Computing Equilibria with Two Known Technologies

Consider a version of the game with two known designs. Since a version of the proposition in subsection 3.2 applies to this game, it is sufficient to compute strategies and values for a *finite* number of values of d (given ϕ). With finite action spaces, an SMPE thus induces a finite partition of the state space into "action regions", indexed by which pair of actions is played, with one region for possible mixed actions. It is clear that by looping over all possible action regions it is possible to find an equilibrium. Of course, in general this is excessively costly. The key to our algorithm is to come up with an initial guess that makes the procedure practical.

- Step 1A: Iterate on the Bellman equations, until further reduction in the difference between iterates appears impossible.
- Step 1B: Use last iterate to construct an initial guess for the partition of the state space into equilibrium regions.
- Step 1C: Select a partition of the state space into action regions (Loop begins with partitions close to the initial guess, than moves further and further away).
- Step 1D: For any guess of action region, solve the pair of second order difference equations *with varying coefficients* for values induced by the Bellman equations, given the guess of action regions. The boundary conditions are $V^k(d^{hi}, \phi) = \frac{R^k(d^{hi})}{1-\beta}$ and $V^k(d^{lo}, \phi) = \frac{R^k(d^{lo})}{1-\beta}$.
- Check whether the values computed in 1D are consistent with the guess from 1C being an equilibrium. If not, repeat 1C.

Step 2: Backtracking the Learning Process

Let $\Omega_T = \{(d', S', T', \phi') \in \Omega : T' = T\}$. Every state $\omega = (d, S, T, \phi)$ has only two possible "successor states" in the set Ω_T, ω itself and $\omega' = (d+1, S, T, \phi)$. For every ω, the expected continuation utility on the right hand side of the Bellman equation depends only on $V^k(\omega)$, $V^k(\omega')$ and values of states in Ω_{T+1} and Ω_{T+2}. Moreover, we know from the proposition that for d large enough in absolute value, $V^k(d, S, T, \phi) = \frac{R^k(d)}{1-\beta}$ for all S,T and ϕ.

- Step 2A: Begin with initial value functions on the sets Ω_{T+1} and Ω_{T+2}, computed in the previous step.
- Step 2B: Pick a $\hat{d}$ high enough such that $\hat{d} > d^{hi}$
- Step 2C: Solve the Bellman equations for states of the form $(\hat{d}-1, S, T, \phi)$ as non-linear equations in the values $V^k(\hat{d}-1, S, T, \phi)$; $k = I, S$.[13] Solving these nonlinear equations amounts to *simultaneously* determining the continuation values and the equilibrium strategies in all states of the form $(\hat{d}-1, S, T, \phi)$

[13] There are two equations for every pair (S, ϕ), where $S \in \{0, 1, ..T\}$ and $\phi \in \{0, 1\}$.

- Step 2D: Repeat procedure for states of the form $(\hat{d}-2, S, T, \phi)$ and so on, backtracking through all $d < d^{hi}$ to find the values and strategies for all states in Ω_T
- Reduce T, repeat step 2B.

References

[1] Aghion, P., Espinosa, M. P. and Jullien, B. (1993), "Dynamic Duopoly with Learning through Market Experimentation" *Economic Theory*, Vol. 3, pp. 517-39.

[2] Aghion P., Harris C., Howitt P. and Vickers J. (2001), "Competition, Imitation and Growth with Step-by-Step Innovation", *Review of Economic Studies*, vol. 68, pp. 467-492.

[3] Anderson S. P., De Palma A. and Thiesse J.F. (1992), *Discrete Choice Theory of Product Differentiation*, MIT Press, Cambridge, Ma.

[4] Bergemann, D. and Valimaki, J. (1996) "Learning and Strategic Pricing" *Econometrica*, Vol.64, 1125:59.

[5] Besanko, D. and U. Doraszelski (2004), "Capacity Dynamics and Endogenous Asymmetries in Firm Size", *Rand Journal of Economics*, Vol. 35, No. 1, pp. 23-49.

[6] Bhattacharya S., Chatterjee K. and Samuelson L. W. (1986), "Sequential Research and the Adoption of Innovations", *Oxford Economic Papers*, vol. 38, pp. 219-243.

[7] Bolton, P. and Harris, C. (1999) "Strategic Experimentation", *Econometrica*, Vol. 67, 349-374.

[8] Budd C., Harris C. and Vickers J. (1993), "A model of the Evolution of Duopoly: Does the Asymmetry between Firms Tend to Increase or Decrease?", *Review of Economic Studies*, vol. 60, pp. 543-573.

[9] Chari V. V. and Hopenhayn H. (1991), "Vintage Human Capital, Growth and the Diffusion of New Technology", *Journal of Political Economy*, vol. 99, pp. 1142-1165.

[10] Christensen C. M. (1993), "The Rigid Disk Drive Industry: A History of Commercial and Technological Turbulence", *Business History Review,* vol. 67, pp. 531-588.

[11] Christensen C. M. (1997), *The Innovator's Dilemma*, Harvard Business School Press, Boston.

[12] Doraszelski, U. and Markovich, S. (2004) "Advertising Dynamics and Competitive Advantage", mimeo, Northwestern University.

[13] Dutta P. and Rustichini A. (1996), "(s,S)-equilibria for Stochastic Games with an Application to Product Innovation", *Journal of Economic Theory*, vol. 97, pp. 1-39

[14] Ericson R. and Pakes A. (1995), "Markov-Perfect Industry Dynamics: A Framework for Empirical Work", *Review of Economic Studies*, vol. 62, pp. 53-82.

[15] Foster R. N. (1986), *Innovation. The Attacker's Advantage*, Summit Books, New York.

[16] Franco, A. and Filson, D. (2002) "Spin-outs: Knowledge Diffusion through Employee Mobility", Federal Reserve Bank of Minneapolis Staff Report n. 272.

[17] Fudenberg, Drew; Gilbert, Richard J.; Stiglitz, Joseph E. and Tirole, Jean (1983) "Preemption, Leapfrogging, and Competition in Patent Races" *European Economic Review* 22:3-31.

[18] Harris C. (1989), *Dynamic Competition for Market Share: An Undiscounted Model*, mimeo, Nuffield College.

[19] Harris C. and Vickers J. (1993), "Racing with Uncertainty", *Review of Economic Studies*, vol. 54, pp. 1-21.

[20] Henderson R. M. and Clark K. B. (1990), "Architectural Innovation: The Reconfiguration of Existing Product Technologies and the Failure of Established Firms", *Administrative Science Quarterly*, vol. 35, pp. 9-30.

[21] Iansiti M. and Khanna T. (1995), "Technological Evolution, System Architecture and the Obsolescence of Firm Capabilities", *Industrial and Corporate Change*, vol. 4, pp. 333-361.

[22] Jensen R. (1982), "Adoption and Diffusion of an Innovation of Uncertain Profitability", *Journal of Economic Theory*, vol. 27, pp. 182-193.

[23] Jovanovic B. and MacDonald G. M. (1994), "Competitive Diffusion", *Journal of Political Economy*, vol. 102, pp. 24-52.

[24] Jovanovic B. and Nyarko Y. (1996), "Learning by doing and the Choice of Technology", *Econometrica*, vol. 64, pp. 1299-1310.

[25] Judd, Kenneth (1985) "Closed Loop Equilibrium in a Multistage Innovation Race", working paper, Northwestern University

[26] Kaplan S, Murray F. and Henderson R. (2003), "Discontinuities and Senior Management: Assessing the Role of Recognition in Pharmaceutical Firm Response to Biotechnology", *Industrial and Corporate Change*, vol. 12, pp. 203-233.

[27] Karshenas M. and Stoneman P. L. (1993), "Rank, Stock. Order, and Epidemic Effects in the Diffusion of New Process Technologies: an Empirical Model", *RAND Journal of Economics*, vol. 24, pp. 503-528.

[28] Maskin E. and Tirole J. (1994), Markov Perfect Equilibrium, Economic Theory Discussion Paper, n. 23, Harvard Institute for Economic Research.

[29] McCardle K. (1985), Information Acquisition and the Adoption of New Technology, *Management Science*, vol. 31, pp. 1372-1389.

[30] Pakes A. and McGuire P. (1994), "Computing Markov-Perfect Nash Equilibria: Numerical Implications of a Dynamic Differentiated Product Model", *RAND Journal of Economics*, vol. 25, pp. 555-589.

[31] Parente S. L. (1994), "Technology Adoption, Learning-by-Doing, and Economic Growth", *Journal of Economic Theory*, vol. 63, pp. 346-369.

[32] Tornell (1997), "Rational Atrophy: the US Steel Industry", NBER working paper n. 6084.

[33] Tushman M. L. and Anderson P. (1990), "Technological Discontinuities and Dominant Designs: A Cyclical Model of Technological Change", *Administrative Science Quarterly*, vol. 35, pp. 604-633.

[34] Tushman M. L. and Anderson P. (1986), "Technological Discontinuities and Organizational Environments", *Administrative Science Quarterly*, vol. 31, pp. 439-465.

[35] Zucker L. G. and Darby R. (1996), "Costly Information: Firm Transformations, Exit, or Persistent Failure", *American Behavioral Scientist*, vol. 39, pp. 959-974.

Figure 1: Revenue function of the within-period pricing game

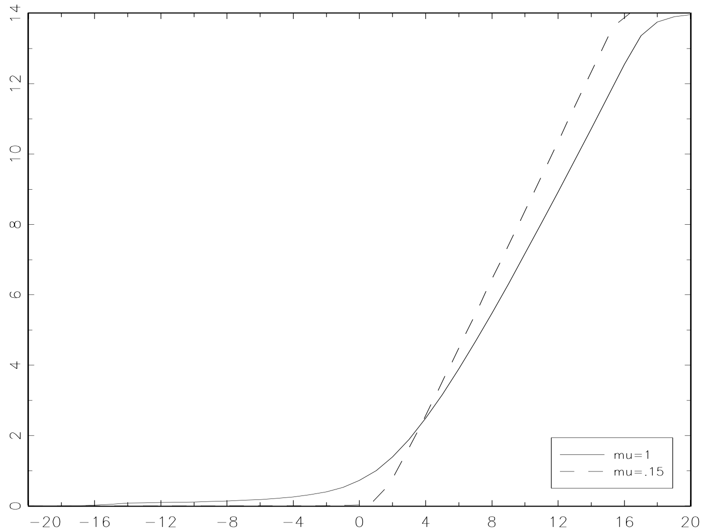

The horizontal axis measures the difference in performance. The vertical axis measures per period profits. The parameter μ governs the substitutability of products of different performance (cf. Appendix A.1). The other parameters are set to their baseline values: the marginal cost of production is $\bar{c} = 1$ and the maximum amount a consumer is willing to spend for a unit of the good is y=15.

Figure 2: Equilibrium value function and equilibrium regions for the one technology certainty game

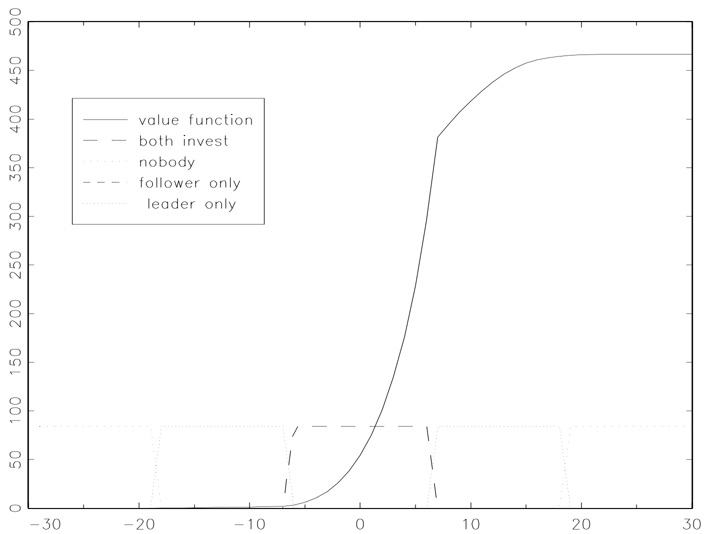

The horizontal axis measures the difference in performance; the vertical axis measures the firm's corresponding value function. The different broken lines indicate equilibrium action regions. Parameter values: probability of success q_0=.7, cost of investment in performance accumulation c=.5, $\bar{c} = 1$, y=15, $\mu = 1$.

Figure 3: Certainty-One-Technology-Game: Comparative Statics

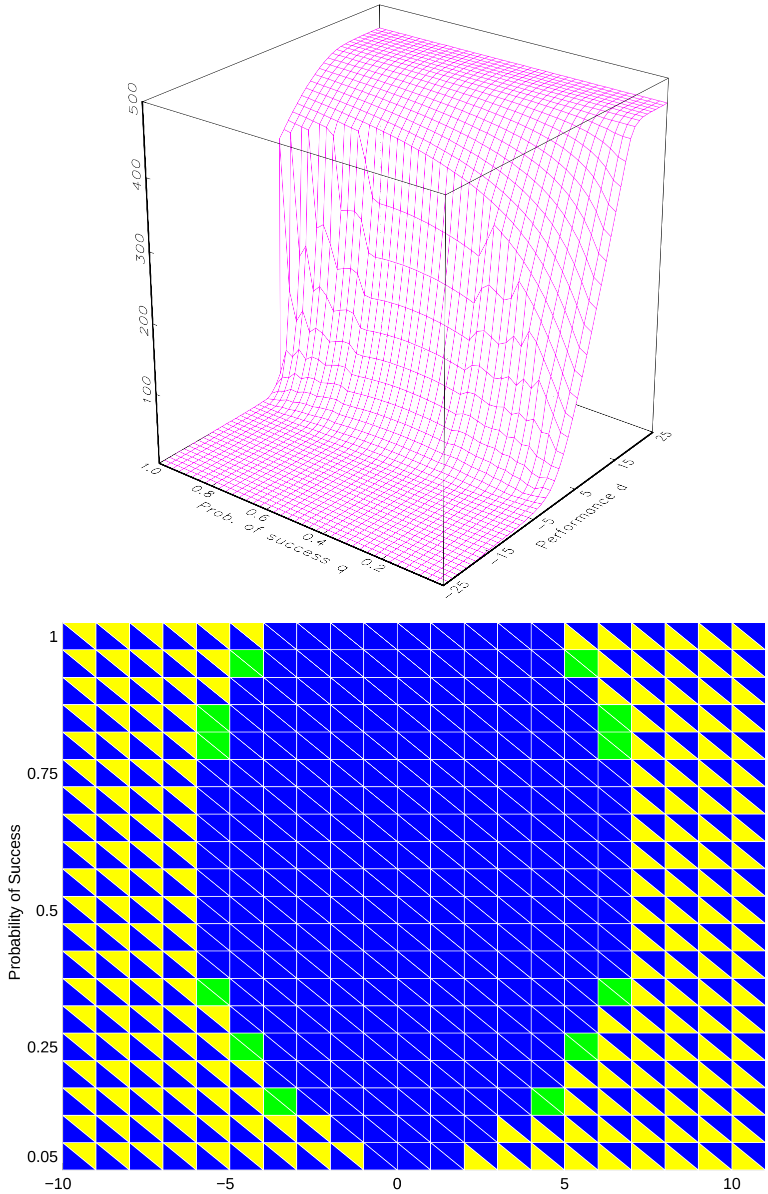

The top panel plots value functions for “certainty, one technology” games that differ by the probability of success under the one, known, design. The only argument of these value functions is the performance advantage.

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 for investment and light/yellow for no investment. Gray/green rectangles indicate mixed actions.

Figure 4: Equilibrium structure: certainty and uncertainty one-technology games

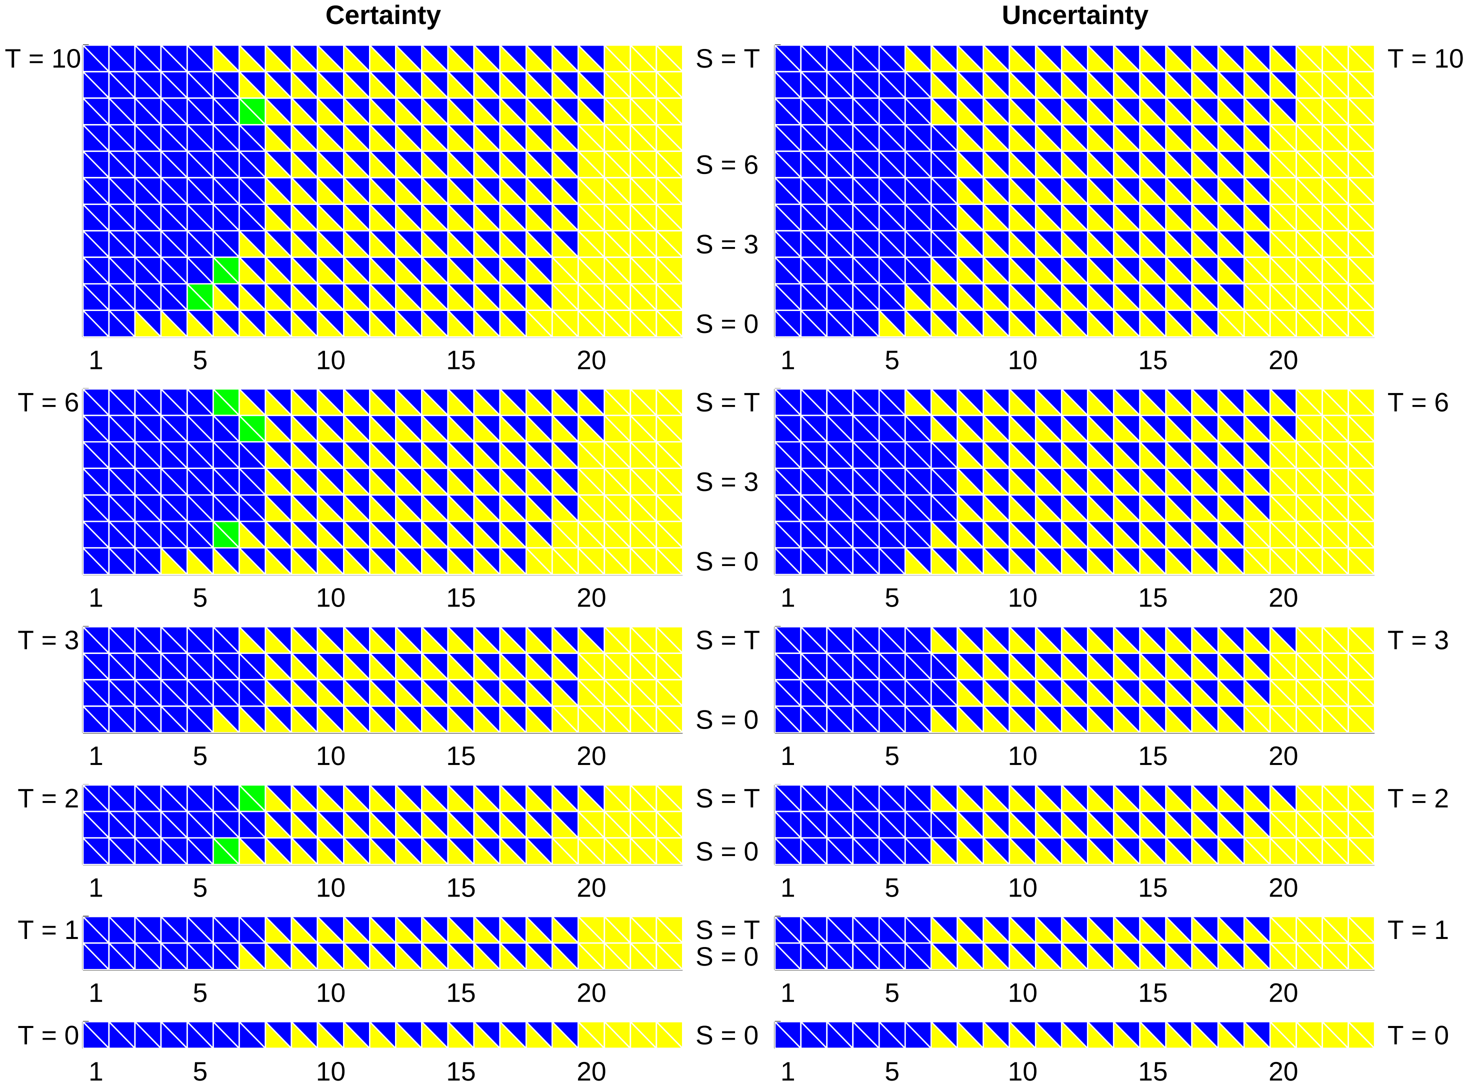

The right-hand set of panels describes equilibrium actions in the uncertainty-one-technology game as a function of the state variables, the performance advantage of the race leader and the beliefs about the new design. The horizontal axis measures the performance advantage of the leader. Each individual 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 for investment in the old technology and light/yellow for no investment. Gray/green rectangles indicate mixed actions.

Figure 5: Equilibrium structure: full model

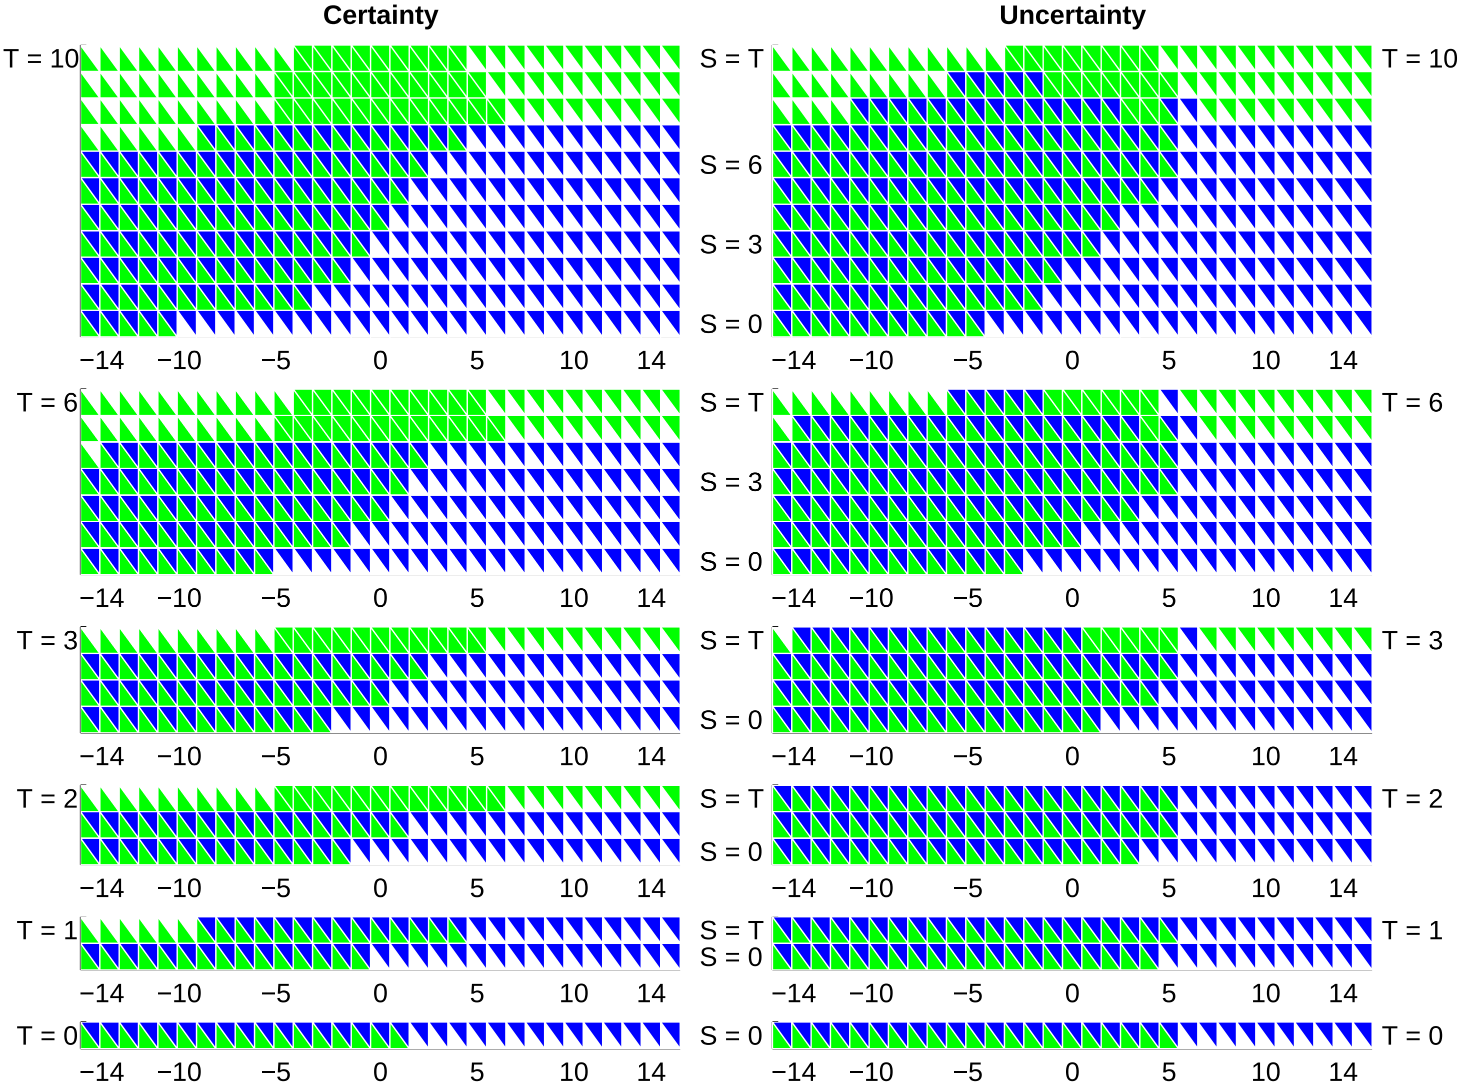

The right-hand set of panels describes equilibrium actions in the full model, with uncertainty about the new design, as a function of the state variables, the performance advantage of the incumbent and the beliefs about the new design. The horizontal axis measures the performance advantage of the incumbent, $d^I - d^S$. 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 for investment in the new technology, dark/blue 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 .7.

Figure 6: The early history before a type II mistake

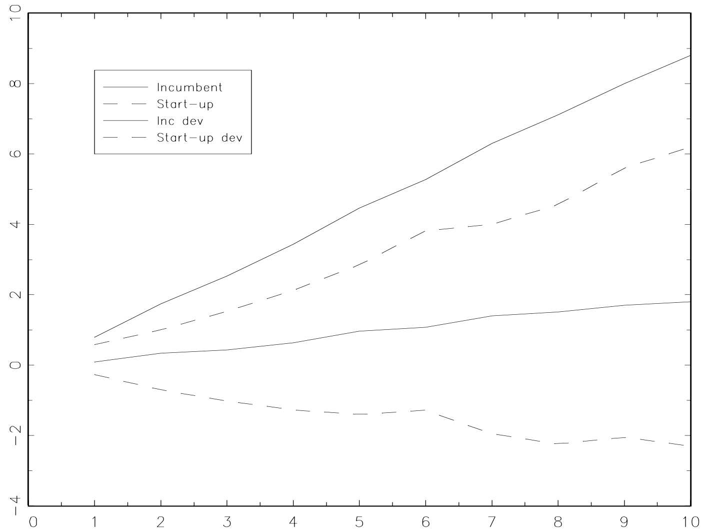

The horizontal axis represents the number of periods for which the game has been played. The top two lines describe, for each firm, the evolution of its performance level (d^I or d^S) along the average path that ends in a type II mistake. The bottom two lines describe, for each firm, the difference between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

Figure 7: The pre-adoption early history and changes in market structure

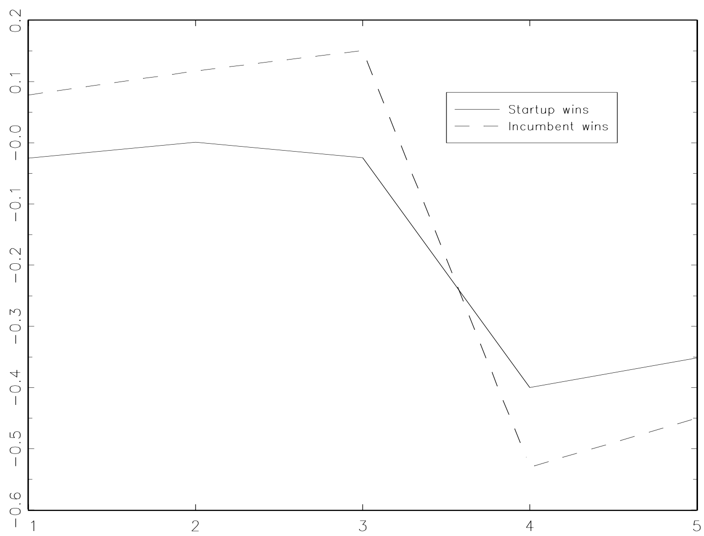

The horizontal axis represents the number of periods T for which the game has been played. A point on the solid line is a pair $(T, d_T^S - \bar{d}_T^S)$. Here d_T^S is the startup's average performance level in period T on paths along which the startup eventually wins, but where the incumbent has not yet adopted by period T. Also, $\bar{d}_T^S$ is the startup's (unconditional) average performance level. The dashed line is constructed the same way, except that d_T^S is obtained by averaging over paths along which the incumbent eventually wins. Parameter values: baseline case with $p = .7, q_0 = .85$, initial advantage $d_0 = 3$

Figure 8: Predicting a change in market structure

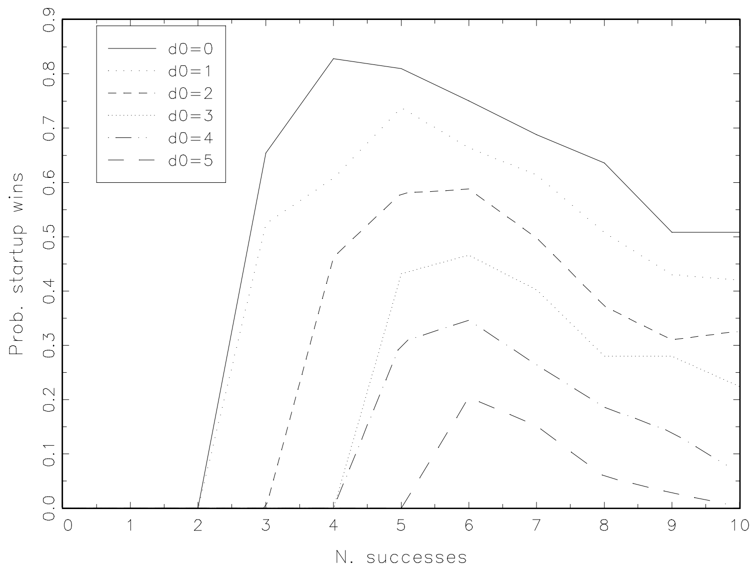

The vertical axis measures the startup's probability of winning, conditional on the state of the game (d, S, T). The number of period of play is fixed at $T = 10$. The number of successes S is measured along the horizontal axis. Different lines correspond to different values of the incumbent's performance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

Figure 9: Disruptive impact of hydraulics in the mechanical excavator market

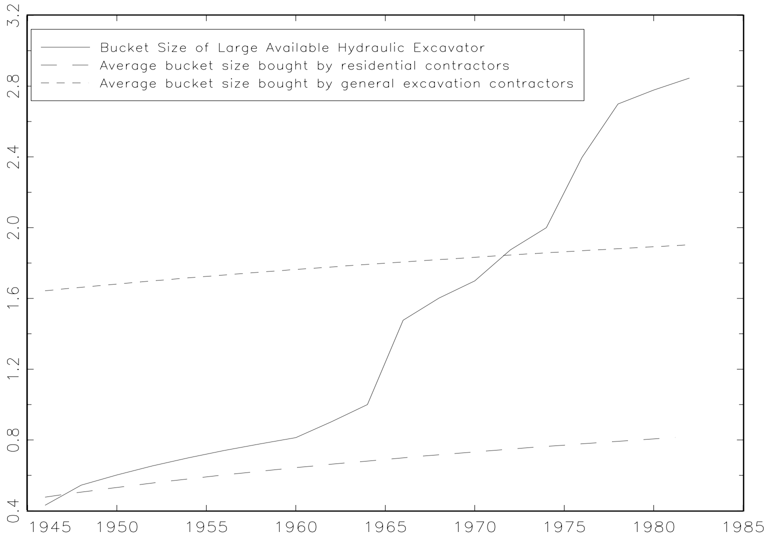

Source: Christensen (1997), pp. 66. Based on data from the Historical Construction Equipment Association. Log scale.

Table 1: Simulation results: summary statistics

d_0	q_0 =.325			q_0 =.5			q_0 =.575			q_0 =.7			q_0 =.775			q_0 =.85			q_0 =.925		
$p = .5$																					
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
$p = .7$																					
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.

Table 2: Average Adoption Time and Advantage at the Time of Adoption, Baseline Case

d_0	q_0 =.575		q_0 =.85		q_0 =.85	
	$p = .5$				$p = .7$	
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.