

Learning-by-Doing, Organizational Forgetting, and Industry Dynamics

David Besanko* Ulrich Doraszelski† Yaroslav Kryukov‡ Mark Satterthwaite§

September 4, 2004

Please do not distribute without permission of the authors

Abstract

It is generally believed that learning-by-doing acts a force that leads to market dominance. Because organizational forgetting can erode a learning-based cost advantage, it might be expected that organizational forgetting would make it less likely that learning-by-doing would lead to market dominance. In this paper we show that this traditional intuition may be incorrect. We analyze a fully dynamic model of price competition when firms face a learning curve and the possibility of organizational forgetting. We show that even though the leader firm may under-price the follower and this price difference may grow as the leader's cost advantage widens, the market may remain unconcentrated in both the short run and long run. And even when learning-by-doing does give rise to a non-trivial degree of market concentration, a steepening of the learning curve does not necessarily translate into a higher degree of concentration. We also show that organizational forgetting does not act as a simple offset to the effects of learning-by-doing. Rather, over an interesting range of parameters, organizational forgetting intensifies pricing rivalry and leads to a greater degree of market concentration. In other words, instead of serving as an antidote to market dominance, organizational forgetting makes a learning-based cost advantage more sustainable and thus makes it more likely that the market will be dominated by a single firm. By extending the model to include entry and exit, we show that predatory pricing can arise endogenously and that organizational forgetting makes predatory behavior more likely to occur. We develop these insights by employing the framework in Ericson and Pakes (1995) to numerically analyze the Markov perfect equilibria (MPE) in a pricing game in a differentiated products duopoly market. A striking feature of our analysis is that in contrast to recent papers that have employed this computational framework, we show that there can be multiple symmetric MPE.

*Kellogg School of Management, Northwestern University, Evanston, IL 60208-2009, U.S.A., d-besanko@kellogg.northwestern.edu.

†Hoover Institution, Stanford University, Stanford, CA 94305-6010, U.S.A., doraszelski@hoover.stanford.edu.

‡Department of Economics, Northwestern University, Evanston, IL 60208-2009, U.S.A., kryukov@northwestern.edu.

§Kellogg School of Management, Northwestern University, Evanston, IL 60208-2009, U.S.A., m-satterthwaite@kellogg.northwestern.edu.

1 Introduction

Empirical studies provide ample support for the hypothesis that learning-by-doing can be an important source of cost advantage. Learning-based cost reductions have been documented in a wide variety of industrial settings including airframes (Alchian 1963), chemical processing (Lieberman 1984), construction of nuclear power plants (Zimmerman 1982), semiconductors (Irwin and Klenow 1994), and shipbuilding (Thompson 2001, Thornton and Thompson 2001). Empirical work has also begun to suggest the possibility that learning-by-doing can be “undone” by organizational forgetting, i.e., the depreciation of the know-how that the firm has built up in the production process. Empirical evidence on organizational forgetting has been presented by Argote and Epple (1990) and Thompson (2003) for the case of World War II Liberty Ships, by Benkard (2000) in the production of wide-body airframes, and by Darr, Argote, and Epple (1995) in the operation of pizza franchises.

It is generally believed that learning-by-doing acts a force that leads to market dominance. The idea is straightforward: a firm with a learning-based cost advantage can profitably under-price its rivals for current sales, moving it further down the learning curve than its rivals. As the firm’s cost advantage widens, it has even more “leg room” to under-price its rivals in the future, eventually allowing the firm to win nearly all sales. Since organizational forgetting can erode a learning-based cost advantage, one expects that organizational forgetting makes it less likely that learning-by-doing leads to market dominance. In other words, if learning-by-doing gives rise to a “winner-take-all” market, organizational forgetting acts as antidote to market dominance by making a firm’s learning-based cost advantage less sustainable.

In this paper we show that this traditional intuition may be incorrect. We analyze a fully dynamic model of price competition with differentiated products when firms face a learning curve and the possibility of organizational forgetting. We show that even though the leader may under-price the follower — a phenomenon Cabral and Riordan (1994) refer to as increasing dominance or ID — and the price difference may grow as the leader’s cost advantage widens (increasing increasing dominance or IID), the market may remain unconcentrated in both the short run and long run. And even when learning-by-doing does give rise to a high degree of market concentration, a steepening of the learning curve does not necessarily translate into a higher degree of concentration. Contrary to the intuition above, we also show that organizational forgetting does not simply offset learning-by-doing. Rather, over an interesting range of parameters, organizational forgetting intensifies pricing rivalry and leads to a greater degree of market concentration. In other words, instead of serving as an antidote to market dominance, organizational forgetting can make a learning-based cost advantage even more sustainable and thus makes market dominance more likely. This result is particularly striking because in our model we employ a specification in which the likelihood of forgetting increases in the amount of know-how that the firm has learned, and so one suspects that a learning-based cost advantage would be especially fragile.¹

We develop these insights by employing the framework in Ericson and Pakes (1995) to numerically solve

¹As we discuss below, this specification is consistent with empirical evidence on learning and forgetting in industrial contexts.

the Markov perfect equilibria (MPE). A striking feature of our analysis is that in contrast to recent papers that have employed this computational framework, we show that there can be multiple symmetric MPE. These multiple equilibria are not an artifact of the mathematical properties of the demand and cost functions we use. Rather, they are grounded in the dynamics of the model and can be interpreted as a bootstrapping phenomenon spawned by organizational forgetting. For a given set of parameter values, different equilibria can give rise to very different industry structures in the short run and the long run. A given set of parameter values may give rise to an equilibrium pricing function in which the leader charges a price that is lower, but not too much lower, than the follower's, allowing both firms to move down their learning curves in tandem. These same parameter values may give rise to a different equilibrium pricing function in which a firm that gains a cost advantage will charge extremely low prices in order to aggressively defend its advantage if it is in immediate danger of losing its cost advantage. Indeed, in some equilibria, a firm that has a significant knowledge lead will launch a price war to prevent the follower firm from moving into a position in which it *might* eventually pose a threat to the leader's cost advantage. In these equilibria, the resulting industry structure can be quite asymmetric. The distinctive role played by organizational forgetting is reflected in the fact that these multiple equilibria only arise in the presence of organizational forgetting. With no organizational forgetting, we find a unique equilibrium in which each firm eventually reaches the bottom of its learning curve.

A number of theoretical studies have explored the competitive implications of the learning curve, primarily focusing on quantity-setting games in finite-horizon models, including Spence (1981), Fudenberg and Tirole (1983), Ghemawat and Spence (1985), Ross (1986), and Cabral and Riordan (1997). The papers that are closest in spirit to ours are Cabral and Riordan (1994), Lewis and Yildirim (2002), and Benkard (2003). Cabral and Riordan (1994) analyze a symmetric MPE in a differentiated-product duopoly market in which firms compete via prices to win a sale in each period. Cabral and Riordan focus on the circumstances under which equilibrium prices exhibit ID and IID. Our analysis differs from Cabral and Riordan in that we consider a model with organizational forgetting and allow for the possibility that neither firm may make a sale in a particular period. Adding organizational forgetting makes it impossible to solve our model analytically, and so unlike Cabral and Riordan, we rely on a computational rather than an analytical approach to the characterization of the MPE. Because we compute the equilibrium pricing function, we can use the theory of Markov processes to determine the transient and ergodic distributions over states, which enables us to develop a characterization of industry dynamics that is much richer than would be possible through a purely analytical approach. A key insight from this analysis is, as noted above, that ID and IID are not sufficient for an industry to be dominated by a market leader.²

Lewis and Yildirim (2002) analyze a MPE between suppliers facing a learning curve, but they focus on the issue of how a single buyer (e.g., a government agency) can optimally design a multi-period procurement auction. This emphasis is very different from the objective of our paper which is to analyze competition

²Nor, as we show, are ID and IID *necessary* for market dominance.

between firms facing a learning curve in an unregulated market in which buyers cannot affect the dynamics of the industry.

Like our paper, Benkard (2003) uses numerical methods to compute the MPE in a model of dynamic competition that includes learning-by-doing, organizational forgetting, and entry and exit. Benkard’s objective is to calibrate his model to the wide-body jet market of the 1970s and 1980s, and he shows that the implied equilibrium dynamics closely track the observed dynamics of price competition in that industry. In contrast to Benkard’s paper, our focus is not on calibrating our model to a particular industry setting. Rather, we show how underlying economic fundamentals can shape industry dynamics, with a particular focus on the role of organizational forgetting.

The organization of the remainder of the paper is as follows. Section 2 describes the model and the approach used in our computations. Section 3 provides an overview of the equilibrium correspondence. Section 4 analyzes pricing behavior in equilibrium, and Section 5 characterizes the industry dynamics implied by it. Section 6 undertakes a number of robustness checks. Section 7 extends the model to include entry and exit and applies the insights of the model to analyze predatory pricing and limit pricing. Section 8 summarizes and concludes.

2 Model

For expositional simplicity, we focus on an industry with two firms without entry and exit. The general model is outlined in Appendix B.

2.1 Firms and states

We consider an infinite-horizon dynamic game in an industry that consists of two firms, indexed by $n \in \{1, 2\}$. Firm n is described by its state $e_n \in \{1, \dots, M\} \equiv \Xi$, where a state describes a firm’s stock of know-how. At any point in time, the industry is characterized by a vector of firms’ states $\mathbf{e} = (e_1, e_2) \in \Xi^2$, where Ξ^2 is the state space. The notation $\mathbf{e}^{[2]}$ denotes the state (e_2, e_1) found by switching the know-how levels of the two firms.

The marginal cost $c(e_n)$ of firm n depends on its accumulated know-how. The firms face a learning curve given by

$$c(e_n) = \begin{cases} \kappa e_n^\eta & \text{if } 1 \leq e_n < m, \\ \kappa m^\eta & \text{if } m \leq e_n \leq M, \end{cases}$$

where $\eta = \frac{\ln \rho}{\ln 2} = \log_2 \rho$ for a learning curve with a slope of ρ percent. Thus, unit cost decreases by $1 - \rho$ percent whenever cumulative experience doubles. Coefficient κ is the marginal cost with minimal experience (normalized to be $e = 1$), and $m < M$ represents the experience level at which the firm reaches the bottom of its experience curve.

Following Cabral and Riordan (1994), we take a period to be a length of time that is just long enough for

at most one firm to make a sale. By making a sale, a firm can add to its stock of know-how. In contrast to Cabral and Riordan (1994), however, we also incorporate organizational forgetting in our model; see Argote and Epple (1990), Darr, Argote, and Epple (1995), Benkard (2000), and Thompson (2003) for empirical evidence. Accordingly, we assume that the evolution of the stock of know-how of firm n is governed the following law of motion:

$$e'_n = e_n + \tilde{q}_n - \tilde{f}_n,$$

where e_n is the firm's know-how in the current period, e'_n is its know-how in the next period, $\tilde{q}_n \in \{0, 1\}$ is the firm's quantity in the current period, and $\tilde{f}_n \in \{0, 1\}$ represents organizational forgetting. If $\tilde{q}_n = 1$, the firm augments its know-how through learning-by-doing, while if $\tilde{f}_n = 1$, the firm loses a unit of know-how through organizational forgetting.

We let $\delta(e_n) = \Pr(\tilde{f}_n = 1)$ denote the probability that firm n 's know-how depreciates in the current period, and we assume that this probability is increasing in e_n . Assuming that the probability of forgetting increases in the amount of know-how has several advantages. First, experimental evidence on forgetting in the industrial psychology literature suggests that the rate of forgetting is an increasing function of the amount learned (Bailey 1989). Second, by assuming that $\delta(e_n)$ is increasing, it can be shown that the expected decay of know-how in the absence of future learning is a convex function of time.³ This phenomenon (known in the psychology literature as Jost's Second Law) is consistent with experimental evidence on forgetting by individuals (Wixted and Ebbesen 1991). Finally, a specification in which $\delta(e_n)$ increases in e_n is conceptually similar to the "capital-stock" model employed in empirical work on organizational forgetting. In the capital-stock model, the depreciation of know-how is proportional to the existing stock of know-how, and so to counteract this, the accumulation of new knowledge through learning-by-doing must increase in proportion to the stock of existing know-how.⁴ Our specification has the same property: as the firm's stock of know-how goes up, its probability of winning a sale must increase in order to counteract the depreciation of its know-how.⁵ In the computations described in the next section, we employ the functional form, $\delta(e_n) = 1 - (1 - \delta)^{e_n}$, where $\delta \in [0, 1]$. When $\delta > 0$, this function is strictly increasing and concave in e_n . If $\delta = 0$, then $\delta(e_n) = 0$ for all e_n , and the firm never forgets.⁶

Conditional on firm n making a sale in the current period (an event denoted by w), its know-how changes

³A proof is available on request. If $\delta(e_n)$ was constant in e_n , the decay of know-how would be a linear function of time.

⁴More specifically, the law of motion in a capital stock model of organizational forgetting is $e'_n = (1 - \psi)e_n + q_n$, where ψ is the depreciation rate.

⁵We do not employ a capital stock model of depreciation in this paper in order to keep the state space integer-valued. See Benkard (2003) for an alternative specification that approximates the capital stock model while maintaining an integer-valued state space.

⁶One way to motivate this functional form would be to imagine that forgetting results when skilled workers leave the firm. In particular, suppose that each additional unit produced by the firm enables one more worker to acquire a special job-related skill that increases his/her productivity, and thus a unit e_n of know-how corresponds to the number of workers in firm n who have acquired the special skill. If one or more of these specially skilled workers leaves the firm in the next period, the firm is assumed to lose, on net, the equivalent of one unit of experience. (Think of new workers being hired to replace those that leave, and assume that there is just enough "organizational memory" so that all but one of the new workers can be trained to replicate the skills of the workers they are replacing.) Under this interpretation, the forgetting parameter δ would represent the rate of labor turnover in a given period, and $\delta(e_n)$ would be the probability that at least one worker from the group of e_n workers who possess the special skill leaves the firm in the next period.

according the transition function

$$\Pr(e'_n|e_n, w) = \begin{cases} 1 - \delta(e_n) & \text{if } e'_n = e_n + 1, \\ \delta(e_n) & \text{if } e'_n = e_n. \end{cases}$$

Conditional on firm n not making a sale (an event denoted by l), its know-how changes according to the transition function

$$\Pr(e'_n|e_n, l) = \begin{cases} 1 - \delta(e_n) & \text{if } e'_n = e_n, \\ \delta(e'_n) & \text{if } e'_n = e_n - 1. \end{cases}$$

At the upper and lower boundaries of the state space, we take the transition function to be $\Pr(M|M, w) = 1$ and $\Pr(1|1, l) = 1$, respectively.

2.2 Price competition

The industry draws its customers from a large pool of potential buyers. In each period, one buyer enters the market and makes, at most, one purchase, either from one of the two firms in the industry (“inside” goods 1 or 2) or an alternative product (“outside” good 0) made from a substitute technology.⁷ The net utility that a buyer obtains by purchasing product n is $v_n - p_n + \tilde{\varepsilon}_n$, where v_n is a deterministic component of net utility, and $\tilde{\varepsilon}_n$ is the buyer’s idiosyncratic preference for product n . We assume that the deterministic utility of the inside goods is the same, so $v_1 = v_2 = v$. Further, we assume that the outside good is supplied under conditions of perfect competition that drive its price to its marginal cost c_0 .

A buyer’s idiosyncratic preferences $(\tilde{\varepsilon}_0, \tilde{\varepsilon}_1, \tilde{\varepsilon}_2)$ are unobservable to firms and are assumed to be iid Type 1 extreme value random variables with location parameter 0 and scale parameter σ . The scale parameter represents the extent of horizontal product differentiation, with a lower value of σ corresponding to weaker product differentiation. As $\sigma \rightarrow 0$, the industry becomes a homogeneous product oligopoly.

In any period, the buyer chooses the alternative that gives it the highest net utility. Given the assumed distribution of the idiosyncratic valuations, the probability that firm n makes a sale when prices are $\mathbf{p} \equiv (p_1, p_2)$ is given by the logit specification

$$D_n(\mathbf{p}) = \frac{\exp(\frac{v-p_n}{\sigma})}{\exp(\frac{v_0-c_0}{\sigma}) + \sum_{j=1}^2 \exp(\frac{v-p_j}{\sigma})}.$$

The demand function is symmetric (i.e., $D_1(p_1, p_2) = D_2(p_2, p_1)$) and has three economically meaningful parameters: σ, v , and $v_0 - c_0$. As $v_0 - c_0 \rightarrow -\infty$ (i.e., the inside goods are infinitely better than the outside good), we are in the Cabral and Riordan (1994) setting in which the buyer always purchases from one of the two firms in the industry.

⁷Since there is a different buyer in each period, buyers are non-strategic. Lewis and Yildirim (2002) consider a model of pricing along a learning curve when there is a single buyer who anticipates the dynamics of future competition.

2.3 Bellman equation

Consider an industry that is in state \mathbf{e} . Letting $\beta \in (0, 1)$ denote a firm's discount factor,⁸ the net present value of future cash flows to firm 1 is given by:

$$V_1(\mathbf{e}) = \max_{p_1} \left\{ (p_1 - c(e_1))D_1(p_1, p_2(\mathbf{e})) + \beta \sum_{k=0}^2 D_k(p_1, p_2(\mathbf{e}))\bar{V}_{1k}(\mathbf{e}) \right\}, \quad (1)$$

where $p_2(\mathbf{e})$ denotes the price charged by the rival firm, and

$$\bar{V}_{1k}(\mathbf{e}) \equiv E[V_1(\mathbf{e}') | \mathbf{e}, \text{buyer purchases good } k], \quad k \in \{0, 1, 2\},$$

is the expectation of firm 1's value function, conditional on the buyer purchasing good k in state \mathbf{e} .⁹ Firm 2 has a comparable value function.

Let $\Omega_1(p_1)$ denote the term in brackets in equation (1). Differentiating it with respect to p_1 , and using properties of the logit demand specification¹⁰, we have

$$\frac{\partial \Omega_1}{\partial p_1} = D_1 \left[1 - \frac{(p_1 - c(e_1))}{\sigma} - \frac{\beta}{\sigma} \bar{V}_{11} + \frac{\Omega_1}{\sigma} \right]. \quad (2)$$

Differentiating this again, combining terms, and using the expression in (2) gives us:

$$\frac{\partial^2 \Omega_1}{\partial p_1^2} = -\frac{1}{\sigma} \frac{\partial \Omega_1}{\partial p_1} [1 - 2D_1] - \frac{D_1}{\sigma}.$$

Thus, $\frac{\partial \Omega_1}{\partial p_1} = 0 \Rightarrow \frac{\partial^2 \Omega_1}{\partial p_1^2} = -\frac{D_1}{\sigma} < 0$, i.e., the objective function is strictly quasi-concave and the price choice $p_1(\mathbf{e})$ is therefore unique.

Setting $\frac{\partial \Omega_1}{\partial p_1}$ to 0 and rearranging terms, firm 1's equilibrium price can be shown to solve the following equation:

$$p_1 = c(e_1) + \frac{\sigma}{1 - D_1(p_1, p_2(\mathbf{e}))} - \beta \Phi(p_2(\mathbf{e}), \mathbf{e}), \quad (3)$$

where

$$\Phi(p_2, \mathbf{e}) \equiv \bar{V}_{11}(\mathbf{e}) - [\eta(p_2)\bar{V}_{12}(\mathbf{e}) + (1 - \eta(p_2))\bar{V}_{10}(\mathbf{e})]$$

⁸The discount factor is given by $\beta = \frac{\zeta}{1+r}$, where $r > 0$ denotes the per period discount rate and $1 - \zeta \in [0, 1)$ is an exogenous probability that the industry's order flow vanishes in the next period (e.g., because their products are supplanted by a drastic innovation that draws away its pool of customers).

⁹To illustrate the form of $\bar{V}_{1k}(\mathbf{e})$, consider $\bar{V}_{12}(\mathbf{e})$, the expectation of firm 1's value function given that the buyer purchases from firm 2:

$$\begin{aligned} \bar{V}_{12}(\mathbf{e}) = & \Pr(e_1 - 1 | e_1, l) [\Pr(e_2 | e_2, w) V(e_1 - 1, e_2) + \Pr(e_2 + 1 | e_2, w) V(e_1 - 1, e_2 + 1)] \\ & + \Pr(e_1 | e_1, l) [\Pr(e_2 | e_2, w) V(e_1, e_2) + \Pr(e_2 + 1 | e_2, w) V(e_1, e_2 + 1)]. \end{aligned}$$

¹⁰In particular, we use that $\frac{\partial D_n(\mathbf{p})}{\partial p_n} = -\frac{1}{\sigma} D_n(\mathbf{p})(1 - D_n(\mathbf{p}))$ and $\frac{\partial D_k(\mathbf{p})}{\partial p_n} = \frac{1}{\sigma} D_n(\mathbf{p}) D_k(\mathbf{p})$, $k \neq n$.

and

$$\eta(p_2) \equiv \frac{\exp(\frac{v-p_2}{\sigma})}{\exp(\frac{v-p_2}{\sigma}) + \exp(\frac{v_0-c_0}{\sigma})}.$$

The term $\Phi(p_2, \mathbf{e})$ is firm 1's (undiscounted) "prize" from making the next sale in state \mathbf{e} when firm 2 charges price p_2 . It consists of the difference between the firm's expected value $\bar{V}_{11}(\mathbf{e})$ if it wins the sale and its expected value $\eta(p_2)\bar{V}_{12}(\mathbf{e}) + (1 - \eta(p_2))\bar{V}_{10}(\mathbf{e})$ if it does not win the next sale, where $\eta(p_2)$ is the probability that the rival wins the next sale, conditional on the firm not winning. The prize represents the wedge that makes dynamic pricing behavior differ from static pricing behavior.

2.4 Equilibrium

Because the demand specification is symmetric, we focus attention on symmetric Markov perfect equilibria (MPE). Existence of a symmetric MPE in pure strategies follows from the arguments in Doraszelski and Satterthwaite (2004). In a symmetric MPE, it suffices to determine the value and policy functions of firm 1, and throughout the remainder of the paper, $p^*(\mathbf{e}) \equiv p_1(\mathbf{e})$ denotes the symmetric equilibrium pricing function for firm 1, while $V^*(\mathbf{e}) = V_1(\mathbf{e})$ denotes the corresponding equilibrium value function.¹¹ Further, we let $\bar{V}_k^*(\mathbf{e}) \equiv \bar{V}_{nk}^*(\mathbf{e})$, $k = 0, 1, 2$, denote the expected value functions. Finally, we let $D^*(\mathbf{e}) \equiv D_1(p^*(\mathbf{e}), p^*(\mathbf{e}^{[2]}))$ denote the equilibrium probability that the typical firm makes a sale in state \mathbf{e} .

Given this notation, we can write the symmetric MPE as follows:

$$V^*(\mathbf{e}) = [p^*(\mathbf{e}) - c^*(\mathbf{e})]D^*(\mathbf{e}) + \beta \left[\eta(p^*(\mathbf{e}^{[2]}))\bar{V}_2^*(\mathbf{e}) + (1 - \eta(p^*(\mathbf{e}^{[2]})))\bar{V}_0^*(\mathbf{e}) \right]. \quad (4)$$

$$p^*(\mathbf{e}) = c^*(\mathbf{e}) + \frac{\sigma}{1 - D^*(\mathbf{e})}, \quad (5)$$

where

$$c^*(\mathbf{e}) \equiv c(e_1) - \beta\Phi(p^*(\mathbf{e}^{[2]}), \mathbf{e})$$

is the *virtual marginal cost*. The virtual marginal cost is endogenously determined in equilibrium and equals the actual marginal cost minus the discounted prize from winning. Since the (symmetric) static Nash equilibrium price $p^\dagger(\mathbf{c})$ for marginal costs $\mathbf{c} = (c_1, c_2)$ is given by

$$p^\dagger(\mathbf{c}) = c_1 + \frac{\sigma}{1 - D^\dagger(\mathbf{c})},$$

where $D^\dagger(\mathbf{c}) \equiv D_1(p^\dagger(\mathbf{c}), p^\dagger(\mathbf{c}^{[2]}))$, it follows that the MPE price $p^*(\mathbf{e})$ is the static Nash equilibrium price corresponding to the virtual marginal costs $\mathbf{c}^*(\mathbf{e}) = (c^*(\mathbf{e}), c^*(\mathbf{e}^{[2]}))$; i.e., $p^*(\mathbf{e}) = p^\dagger(\mathbf{c}^*(\mathbf{e}))$.

¹¹We thus obtain the value and pricing functions of firm 2 by switching the arguments of firm 1's pricing and value functions: $p_2(\mathbf{e}) = p_1(\mathbf{e}^{[2]})$. $V_2(\mathbf{e}) = V_1(\mathbf{e}^{[2]})$.

2.5 Computation

To compute a symmetric MPE, we adapt the algorithm described in Pakes and McGuire (1994). The algorithm works iteratively. It takes a value function $\tilde{V}^*(\mathbf{e})$ and pricing function $\tilde{p}^*(\mathbf{e})$ as the starting point for an iteration and generates updated value and policy functions. Each iteration proceeds as follows:

1. First, we use equation (5) to compute an updated pricing policy $p^*(\mathbf{e})$ for firm 1, taking the pricing policies of firm 2 to be $\tilde{p}^*(\mathbf{e}^{[2]})$. In this step, we use $\tilde{V}^*(\mathbf{e})$ to compute the expected value functions, $\bar{V}_0^*(\mathbf{e})$, $\bar{V}_1^*(\mathbf{e})$, and $\bar{V}_2^*(\mathbf{e})$.
2. Next, we use equation (4) to compute an updated value function, $V^*(\mathbf{e})$.
3. The iteration is completed by assigning $V^*(\mathbf{e})$ to $\tilde{V}^*(\mathbf{e})$, and $p^*(\mathbf{e})$ to $\tilde{p}^*(\mathbf{e})$.¹²

The algorithm terminates once the relative change in the value and the policy functions from one iteration to the next is below a pre-specified tolerance. All programs are written in Matlab 6.5. Details are available from the authors upon request.

2.6 Parameterization

In our numerical analysis, will focus on how the learning curve and organizational forgetting influence the shape of the equilibrium pricing function and the industry dynamics that are implied by that pricing behavior. Accordingly, in our baseline analysis, we will report results for a wide range of values of ρ and δ , holding fixed all other parameters. Later, we will check whether the insights from our baseline analysis are robust to changes in the degree of product differentiation, as measured by σ , and the size of the market of the inside goods, as parameterized by $v_0 - c_0$ relative to v . The values of the parameters used in the computations are as follows.

2.6.1 Rate of learning

We compute the equilibrium for the following values of ρ :

$$\{0.95, 0.85, 0.75, 0.65, 0.55, 0.35, 0.15, 0.05\} \equiv \mathbf{R}$$

Empirical estimates of experience curves tend to find slopes in the range of 0.75 to 0.95 (Lieberman 1984, Dutton and Thomas 1984), and so in describing our results, we will often focus on the case $\rho = 0.85$, which can be thought of (loosely) as the “median” slope across a wide array of empirical estimates.

¹²More precisely, we assign a weighted average of $V(\mathbf{e})$ and $\tilde{V}(\mathbf{e})$ to $\tilde{V}(\mathbf{e})$ to help the algorithm converge.

2.6.2 Rate of organizational forgetting

We compute the equilibrium for the following values of δ :

$$\{0, 0.01, 0.02, 0.03, 0.04, 0.05, 0.06, 0.07, 0.08, 0.09, 0.10, 0.30, 0.50, 0.90\} \equiv \mathbf{D}$$

Empirical studies of organizational forgetting that employ a capital stock model have found rates of forgetting that range from 4 percent per month to 25 percent per month. Our model is not strictly comparable to a capital stock model, so it is not immediately obvious what empirically plausible values of δ might be. To shed light on this question, Appendix A develops a “mapping” between the depreciation rate in a capital stock model and the parameter δ in our model. Based on this analysis, it would appear that the empirically relevant range of δ is between 0 and about 0.10. Accordingly, in describing our results, we will focus our attention on cases in this range: $\delta = 0$ (which will serve as a benchmark against which to compare our model to Cabral and Riordan (1994)), $\delta = 0.03$, and $\delta = 0.08$. As we will see below, this set of δ 's generates a rich variety of industry dynamics.

2.6.3 Parameters held fixed for the baseline analysis

- *Attractiveness of the outside alternative, $v_0 - c_0$* : In our baseline analysis, we set $v_0 - c_0 = 0$. In Section 6, we discuss how the equilibrium is affected if we make the outside alternative more attractive ($v_0 - c_0 = 3$) and (infinitely) less attractive ($v_0 - c_0 = -\infty$).
- *Degree of horizontal differentiation, σ* : In our baseline analysis, we set $\sigma = 1$, resulting in a moderately weak degree of horizontal differentiation.¹³ In Section 6, we discuss how the equilibrium changes if horizontal differentiation becomes stronger ($\sigma = 2$) and extremely weak ($\sigma = 0.10$).

2.6.4 Parameters held fixed throughout analysis

- *State space, M* : Throughout the analysis, we consider a state space with 30 possible levels of know-how, i.e., $M = 30$
- *Know-how level at which learning curve flattens out, m* : We assume that the learning curve flattens out at 15 units of know-how, i.e., $m = 15$.
- *Marginal cost at the “top” of the learning curve, κ* : We consider a learning curve with an initial level of marginal cost equal to 10, i.e., $\kappa = 10$. Note that with $\kappa = 10$, along an 85 percent learning curve ($\rho = 0.85$), marginal cost falls from \$10 per unit to approximately \$5.30 per unit at $e = 15$.

¹³To illustrate, if both firms set a price equal to their initial marginal cost $\kappa = 10$, the own price elasticity of demand would equal -6.67. As firms drop their prices in tandem, the own-price elasticity of demand falls, reaching a level of -3.28 if both firms set a price equal to the marginal cost at the bottom of an 85 percent learning curve.

- *Deterministic component of utility of the inside goods, v* : Throughout the analysis, we fix $v = 10$. Because $v - \kappa = 0$, note that in the baseline case of $v_0 - c_0 = 0$, when a firm is at the “top” of the learning curve, its product is “on a par” with the outside alternative.
- *Discount factor, β* : We set $\beta = 0.9524$. This corresponds to a variety of scenarios that differ according to the length of a period. For example, it corresponds to a case in which a period is one year, the industry is certain to survive, and the discount rate is 5 percent ($r = 0.05$). It could also correspond to a setting in which a period is one month, the market has a 96.5 percent chance of surviving from one month to the next, and the monthly discount rate is about 1.06 percent (which corresponds to an annual discount rate of about 12.7 percent).¹⁴

3 Equilibrium correspondence

To characterize how the nature of learning-by-doing and organizational forgetting shapes equilibrium pricing behavior and industry dynamics, we compute the MPE for all (ρ, δ) combinations in $\mathbf{P} \times \mathbf{D}$, with all other parameters set to their baseline values. We then use the computed pricing function to construct the Markov transition probability of next period’s state \mathbf{e}' given the current state \mathbf{e} . This allows us to use tools from stochastic process theory to analyze the Markov process of industry dynamics rather than relying on simulations. In particular, we can use the transition probabilities to compute the limiting distribution over states, $\pi^\infty(\mathbf{e})$ implied by the equilibrium.

Table 1 provides an overview of the MPE by showing the expected industry Herfindahl index corresponding to the limiting distribution:

$$H^\infty \equiv \sum_{\tilde{\mathbf{e}}_1} \sum_{\tilde{\mathbf{e}}_2} \left[\frac{D^*(\tilde{\mathbf{e}})^2 + D^*(\tilde{\mathbf{e}}^{[2]})^2}{(D^*(\tilde{\mathbf{e}}) + D^*(\tilde{\mathbf{e}}^{[2]}))^2} \right] \pi^\infty(\tilde{\mathbf{e}}).$$

The Herfindahl index is based on the firms’ conditional market shares, and so it is bounded below by 0.5 and above by 1.

Two points emerge from this overview. First, there may be multiple symmetric MPE for a given set of parameter values. Second, though the equilibria can differ greatly across parameter values, equilibrium pricing behavior can generally be classified into one of three categories that differ in interesting ways. We discuss each point in turn.

¹⁴To put this second scenario in perspective, technology companies such as IBM, Cisco Systems, Microsoft, Intel, Dell, and Sun Microsystems had costs of capital in the range of 11 to 15 percent in the late 1990s. Further, an industry with a survival probability of 96.5 percent from one month to the next has an expected life of 26-27 months. Thus, this second scenario might be thought of as broadly representative of a technology company that receives orders from customers on a once-a-month basis for a product in which the pace of innovative activity in the industry would be expected to make it obsolete in roughly two years.

ρ	0.95		0.85		0.75		0.65		0.55		0.35		0.15		0.05	
δ	F	T	F	T	F	T	F	T	F	T	F	T	F	T	F	T
0	.500	—	.500	—	.500	—	.500	—	.500	—	.500	—	.500	—	.500	—
0.01	.500	—	.500	—	.500	—	.500	—	.500	—	.500	—	.500	—	.500	—
0.02	.500	—	.500	—	.500	—	.500	—	.500	—	.500	—	.500	—	.500	—
0.03	.500	.516	—	.520	—	.521	—	.521	—	.521	—	.521	—	.500	.521	—
0.04	.502	.513	—	.519	—	.521	—	.521	—	.501	.521	—	.523	—	.500	.523
0.05	.504	—	.530	—	.538	—	.580	—	—	.542	—	.527	—	.500	.527	—
0.06	.507	—	.559	—	—	.882	—	.923	—	—	.940	.531	—	.500	.530	—
0.07	.511	—	—	.744	—	.884	—	.921	—	—	.938	—	.954	.500	.534	—
0.08	.515	—	—	.770	—	.883	—	.920	—	—	.937	—	.952	.501	.530	.961
0.09	.519	—	—	.787	—	.882	—	.918	—	—	.935	—	.950	.502	.530	.960
0.1	.523	—	—	.803	—	.882	—	.916	—	—	.933	—	.949	.505	.534	.958
0.3	.517	—	—	.860	—	.922	—	.936	—	—	.942	—	.945	—	.945	.945
0.5	.508	—	—	.597	—	.897	—	.956	—	—	.965	—	.969	—	.969	.968
0.7	.504	—	—	.539	—	.615	—	.762	—	—	.904	—	.965	—	.975	.977
0.9	.501	—	—	.511	—	.532	—	.561	—	—	.591	—	.672	—	.766	.803

Table 1: Expected Herfindahl index under limiting distribution for $(\rho, \delta) \in \mathbf{R} \times \mathbf{D}$.




Figure 1: Equilibrium correspondence (expected Herfindahl index) for $\rho = 0.85$ and $\delta \in D$.

3.1 Multiple equilibria

Table 1 illustrates that for some parameter combinations we were able to find two, or even three, symmetric MPE. For example, for $\rho = 0.85$, $\delta = 0.03$, we found two equilibria: one in which the long-run expected Herfindahl is 0.500, and another in which the long-run expected Herfindahl is 0.516. The multiple equilibria were discovered by means of a systematic search that employed various starting values for $p^*(\cdot)$ and $V^*(\cdot)$. Specifically, our approach is to start calculation from the equilibrium for a neighboring set of parameter values, and check if the solution is different from the already known equilibrium.

Generically, we would expect that for any particular combination of parameter values, the number of symmetric MPE is odd (Herings and Peeters 2004). This suggests that for the case of $\rho = 0.85$, the equilibrium correspondence has the shape shown in Figure 1, which in turn suggests that in this case there may be as many as three equilibria when $\delta = 0.03$, of which we have computed two, those corresponding to points A or C on the equilibrium correspondence. Because our computational algorithm relies on iterated best responses, the equilibria that we have computed are stable in the sense that if we chose starting values close enough to the equilibria at A or C , the iterated best response algorithm will lead us back to the equilibria at A or C . By contrast, the conjectured equilibrium at B is likely to be unstable in the sense that even slight perturbations away from B will lead us either to A or C . In this respect, the equilibria we have calculated are appealing because one might imagine that a process of iterated best responses would be one

way that firms would work their way toward an equilibrium.

Given the assumed properties of demand and cost, the static Nash equilibrium in our model can be shown to be unique. Further, we find unique symmetric MPE in the absence of organizational forgetting. Thus, the multiple equilibria in our model seem to arise because of the dynamic structure of the model with organizational forgetting, and not because of any mathematical properties that flow from our specification of demand or the learning curve. Following our discussion of equilibrium pricing behavior and the resulting industry dynamics, we will develop an intuition for why the multiple equilibria arise in our model when there is organizational forgetting.

3.2 Classifying equilibria

Table 1 suggests that there are three broad categories of equilibria. First, we have equilibria in which $H^\infty \approx 0.500$ indicating that in the long run we end up with two equal-sized firms. An example is the case where $\rho = 0.85, \delta = 0$ (no organizational forgetting) or the “first” equilibrium when $\rho = 0.85, \delta = 0.03$. As can be seen from the top two panels of Figure 2, except in a neighborhood of $\mathbf{e} = (1, 1)$, a firm’s equilibrium price is not particularly sensitive to its own state or that of its rival. For this reason, we refer to these as *flat equilibria* (labeled F in Table 1). Although a flat equilibrium might exhibit a range of intense price competition in a neighborhood of $\mathbf{e} = (1, 1)$ — what we will refer to as a “well” — price competition in a flat equilibrium is not particularly intense once firms begin to make modest progress in moving down their learning curves.

Second, we have equilibria in which H^∞ is between 0.500 to approximately 0.580, indicating that in the long run we end up with an asymmetric oligopoly in which one firm has a market share between 50 to 70 percent, while the other firm has a market share between 30 and 50 percent. An example is the “second” equilibrium when $\rho = 0.85, \delta = 0.03$. As can be seen from the graph of $p^*(\mathbf{e})$ in the bottom left panel in Figure 2, the equilibrium pricing function not only has a well but it also exhibits a pronounced “trench” running along the diagonal. For this reason, we refer to these as *trenchy equilibria* (labeled T in Table 1). In a trenchy equilibrium, price competition between equally experienced firms is fairly intense. Price competition abates only when firms move into the asymmetric region of the state space.

Finally, we have equilibria in which H^∞ is extremely high (often in the range between about 0.7 and 0.95), indicating that in the long run we end up with a single firm that essentially dominates the market. An example of this is the case where $\rho = 0.85, \delta = 0.08$. As can be seen from the graph of $p^*(\mathbf{e})$ in the bottom right panel of Figure 2, the equilibrium policy function not only has a deep trench running along the diagonal, it also has a “sideways trench” running along the edges of the state space. For this reason, we refer to equilibria of this type as *extra trenchy* (labeled at X in Table 1). In an extra-trenchy equilibrium, price competition between equally experienced firms is extremely intense; furthermore, there are regions in the state space where there is also intense price competition between firm with a significant cost advantage and its disadvantaged rival.

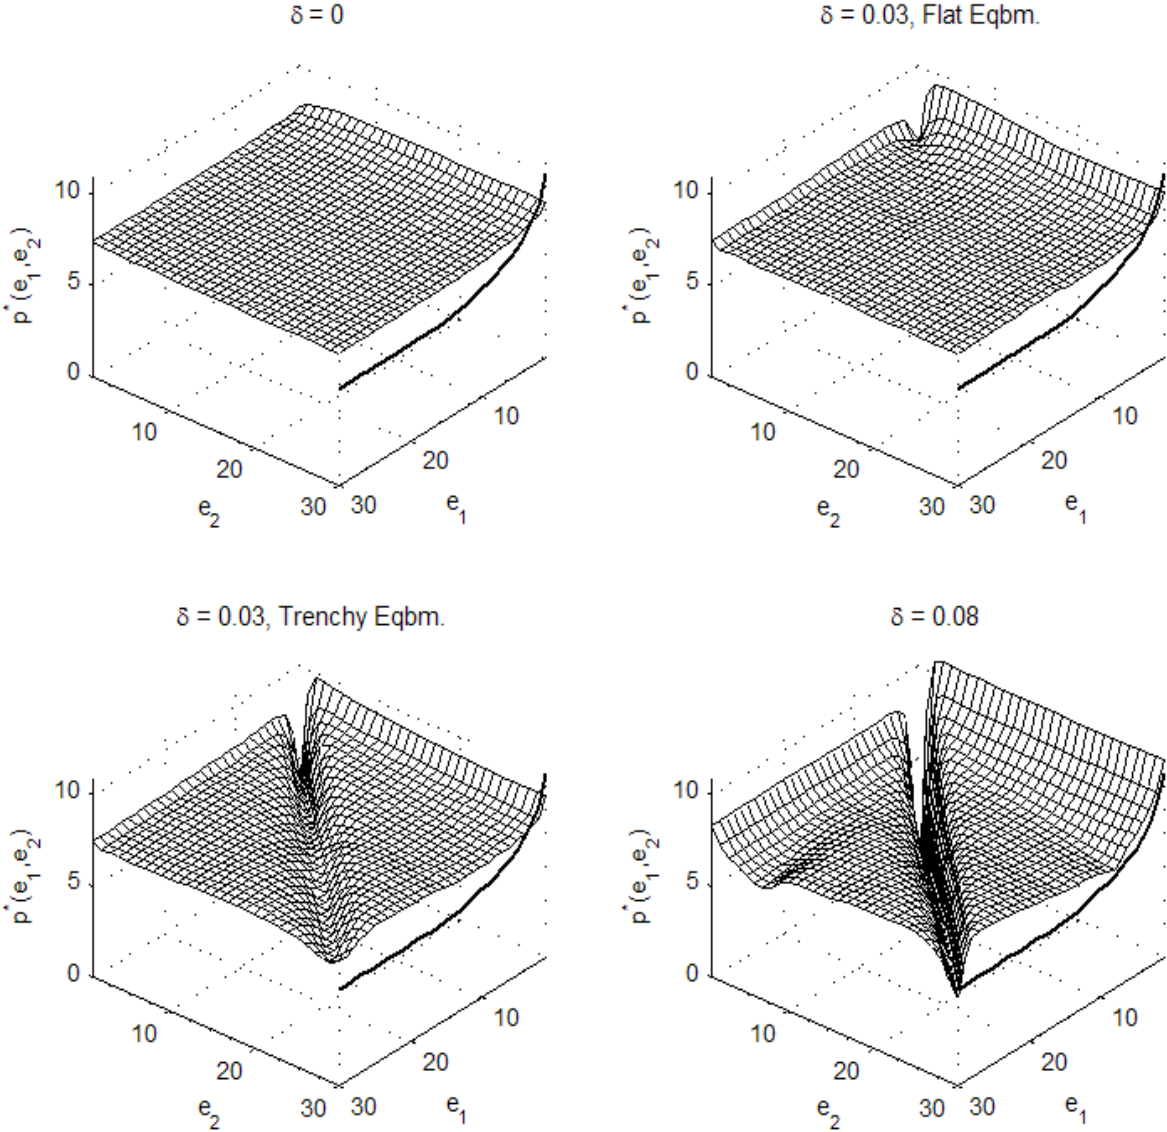


Figure 2: Price function for $\rho = 0.85$ and $\delta = 0, 0.03$ (both equilibria), 0.08 . Line in $e_2 = 30$ plane is the cost function $c(e_1)$.

Figure 3 shows the value functions corresponding to the equilibrium pricing functions in Figure 2. The value functions corresponding to the trenchy and extra-trenchy equilibria show that both a leader and a follower experience a drop in value as the industry moves from an asymmetric state near the diagonal to a state located along the diagonal; in other words, the diagonal trench on the pricing function is mirrored by a trench in the value function. Further, in the value function corresponding to the extra-trenchy equilibrium, the value of being the dominant firm is significant, while the value associated with being the follower is very close to zero. Further, the value associated with the initial state $(1, 1)$ is also low, indicating that firms expect to dissipate significant profitability in the race to become the industry leader.




Figure 3: Value functions for $\rho = 0.85$ and $\delta = 0, 0.03$ (both equilibria), 0.08 .

For each (ρ, δ) pair in Table 1, we have classified the equilibria into one of three categories: F , T , X . Although we motivated the different classes of equilibria by pointing out that they give rise to significantly different expected Herfindahl indexes in the long run, the classification of equilibria in Table 1 is actually based on an algorithm that identifies wells, diagonal trenches, and sideways trenches through direct examination of the equilibrium pricing functions. An equilibrium is classified as flat if the pricing function has no diagonal trench. An equilibrium is classified as trenchy equilibria if it has a diagonal trench, and either no sideways trenches or very narrow and shallow sideways trenches that are close to the axis of state space. An equilibrium is classified as extra-trenchy if it has a diagonal trench and sufficiently deep sideways trenches that are several states wide.

4 Pricing behavior

The three categories of equilibria give rise to different dynamic pricing behavior. In this section, we explore the nature of equilibrium pricing behavior in greater detail.

4.1 Learning-by-doing

It is useful to compare the equilibrium pricing behavior in our dynamic model to the pricing behavior that would arise in the Nash equilibrium in a static model without learning or forgetting. We begin with a definition:

Definition 1 *The value function $V^*(\mathbf{e})$ is locally non-decreasing in its first argument at state \mathbf{e} if $V^*(e_1 + 1, z) \geq V^*(e_1, z) \geq V^*(e_1 - 1, z)$ for $z \in \{e_2 - 1, e_2, e_2 + 1\}$, and it is locally non-increasing in its second argument at state (e_1, e_2) if $V^*(z, e_2 - 1) \geq V^*(z, e_2) \geq V^*(z, e_2 + 1)$ for $z \in \{e_1 - 1, e_1, e_1 + 1\}$. The value function is well-behaved at \mathbf{e} if it is locally non-decreasing in its first argument and locally non-increasing in its second argument at (e_1, e_2) .*

Our computations illustrate that the equilibrium value function is not *guaranteed* to be well behaved, although it *typically is* well behaved over much of the state space. When the value function is well behaved, we can unambiguously compare the MPE prices to the static Nash equilibrium prices:

Proposition 1 *If the value function $V^*(\mathbf{e})$ is well behaved in state \mathbf{e} , then a firm's prize from winning is non-negative, i.e., $\Phi(p, \mathbf{e}) \geq 0$. Furthermore, a firm's virtual marginal cost cannot exceed its actual marginal cost, and its equilibrium price in state \mathbf{e} cannot exceed the static equilibrium price corresponding to the cost pair $\mathbf{c}(\mathbf{e}) \equiv (c(e_1), c(e_2))$, i.e., $\mathbf{c}^*(\mathbf{e}) \leq \mathbf{c}(\mathbf{e})$ and $p^*(\mathbf{e}) \leq p^\dagger(\mathbf{c}(\mathbf{e}))$.*

Proof 1 *Note that*

$$\Phi(p, \mathbf{e}) \equiv \eta(p) \left[\bar{V}_1^*(\mathbf{e}) - \bar{V}_2^*(\mathbf{e}) \right] + (1 - \eta(p)) \left[\bar{V}_1^*(\mathbf{e}) - \bar{V}_0^*(\mathbf{e}) \right].$$

It is straightforward (but tedious) to establish that if $V^(\cdot)$ is well behaved at \mathbf{e} , $\bar{V}_1^*(\mathbf{e}) - \bar{V}_2^*(\mathbf{e}) \geq 0$ and $\bar{V}_1^*(\mathbf{e}) - \bar{V}_0^*(\mathbf{e}) \geq 0$, which implies $\Phi(p, \mathbf{e}) \geq 0$. This, in turn, implies that $c^*(\mathbf{e}) \leq c(e_1)$ and $c^*(\mathbf{e}^{[2]}) \leq c(e_2)$, so that $\mathbf{c}^*(\mathbf{e}) \leq \mathbf{c}(\mathbf{e})$. Because the static Nash equilibrium price $p^\dagger(\mathbf{c})$ is non-decreasing in \mathbf{c} (Vives 1999), it follows that $p^*(\mathbf{e}) = p^\dagger(\mathbf{c}^*(\mathbf{e})) \leq p^\dagger(\mathbf{c}(\mathbf{e}))$.*

This proposition highlights the fundamental economic impact of learning-by-doing. Whenever the value function is well-behaved, a firm anticipates a non-negative prize from winning the next sale. When this prize is strictly positive, the firm acts “as if” its marginal cost is lower than its current out-of-pocket cost $c(e_n)$, inducing each firm, in equilibrium, to price more aggressively than it would have in a static Nash equilibrium. In effect, the amount by which each firm lowers its price below the static Nash equilibrium price is a measure of the extent to which the firm uses price cuts to invest in lowering its future costs.

4.2 Organizational forgetting

Because organizational forgetting counteracts the rate of learning and makes gains in know-how more transitory, one might expect that firms would be more reluctant to invest in the acquisition of know-how through aggressive price competition when there is organizational forgetting. But the diagrams in Figure 2 suggest a different conclusion: organizational forgetting can actually intensify price competition. When we have an 85 percent learning curve with no organizational forgetting ($\delta = 0$), we have a flat equilibrium in which in nearly all states the equilibrium price is very close to the static Nash equilibrium price corresponding to the marginal cost $c(m)$ at the bottom of the learning curve. This reflects the well-understood intuition that when the discount factor β is close to 1, firms facing a deterministic learning curve price “as if” their marginal cost is $c(m)$ (Spence 1981, Cabral and Riordan 1994). But if $\delta = 0.03$, a modest degree of organizational forgetting, we obtain two equilibria, and in each of these equilibria there is fairly intense price competition at the top of the learning curve (in states near $(1, 1)$). Further, in the second of the two equilibria (in the lower right panel of Figure 2), there is also a diagonal trench where equally experienced firms engage in intense price competition. If organizational forgetting becomes even stronger ($\delta = 0.08$), price competition becomes even more intense throughout most of the state space.

The intuitive explanation for the impact of organizational forgetting is this. Organizational forgetting creates a “current” that firms must “fight against” as they attempt to reduce costs by accumulating experience. When there is organizational forgetting, a firm increases its stock of know-how by making sales at a rate that exceeds the rate at which its know-how depreciates through organizational forgetting. When two firms have the same know-how, each firm will set the same price and have an equal probability of making a sale. This probability cannot exceed 0.5, and may be less than 0.5 depending on the attractiveness of the outside alternative. Given this, it may be virtually impossible for two equally experienced firms to move all the way down the learning curve, even though it might be quite possible for one firm to do so. This phenomenon gives rise to two opposing forces. On the one hand, organizational forgetting makes it less attractive for the firms to set low prices when they are at (or close to) the top of the learning curve: after all, why invest in accumulating know-how through price cuts when your know-how gains will be transitory (and when your rival will also be unable to sustain its know-how gains). This is the “investment-stifling” effect of organizational forgetting, and it works to soften price competition in states close to the top of the learning curve. On the other hand, though, because organizational forgetting can give rise to a situation in which the market can support just one low-cost firm, each firm has a strong incentive to move off — and stay off — the diagonal of the state space, even in states at the bottom of the learning curve. We call this the “preemption” effect of organizational forgetting, and it works to intensify price competition.

We can see the trade-off between the investment-stifling and preemption effects in the third and fourth columns of Table 2 which show the equilibrium prices $p^*(1, 1)$ and $p^*(2, 2)$ for symmetric firms at or near the top of an 85 percent learning curve. As δ increases from 0 to about 0.07-0.08, price competition at the top of the learning curve intensifies, indicating that the preemption effect dominates the investment-

stifling effect over this range of δ . However, as δ increases beyond 0.08, price competition at the top of the learning curve begins to wane (although it still remains much more intense than it would be in the absence of organizational forgetting, until we reach the extremely high value of δ of 0.30). Thus, the investment-stifling effect dominates the preemption effect only when δ becomes sufficiently large.

δ	Equilib. type	$p^*(1, 1)$	$p^*(2, 2)$	$p^*(14, 14)$	$p^*(15, 15)$
0	<i>F</i>	7.80	7.78	7.24	7.24
0.01	<i>F</i>	7.64	7.65	7.22	7.23
0.02	<i>F</i>	7.10	7.37	7.19	7.22
0.03	<i>F</i>	4.87	6.25	7.14	7.21
0.03	<i>T</i>	2.89	4.03	5.99	6.13
0.04	<i>F</i>	1.25	3.17	6.80	6.99
0.04	<i>T</i>	1.36	3.03	5.65	5.94
0.05	<i>T</i>	-0.83	1.31	4.84	5.06
0.06	<i>T</i>	-3.63	-1.71	4.41	4.56
0.07	<i>X</i>	-5.08	-3.95	4.02	4.13
0.08	<i>X</i>	-4.59	-4.01	3.05	3.21
0.09	<i>X</i>	-3.94	-3.84	2.10	2.25
0.10	<i>X</i>	-3.19	-3.52	1.25	1.37
0.30	<i>X</i>	8.84	5.60	-5.54	-5.55
0.50	<i>X</i>	10.80	8.82	-5.20	-5.52
0.70	<i>X</i>	11.00	9.24	-5.37	-5.71
0.90	<i>X</i>	11.20	9.37	-5.57	-5.90

Table 2: Prices in symmetric states at the top and the bottom of the learning curve, $\rho = 0.85$. (Equilibrium classifications: *F* = flat equilibrium; *T* = trenchy equilibrium; *X* = extra-trenchy equilibrium.)

The trade-off between the investment stifling and preemption effects is different if one or both firms happen to have reached the bottom, or close to the bottom, of the learning curve. At the bottom of the learning curve, the investment-stifling effect of organizational forgetting is virtually absent: cost reductions from accumulating further know-how are zero or close to zero anyway, so it really does not matter that much that organizational forgetting makes those small gains more transitory. However, the preemption effect continues to operate. We see this in the fourth and fifth columns of Table 2: when firms are at, or near the bottom, of an 85 percent learning curve (states (14, 14) and (15, 15)), increases in δ tend to intensify price competition.

The preemption effect may be so strong that a firm may be willing to set price below its marginal cost *even when it is at the bottom of the learning curve*. A straightforward extension of Theorem 4.4 in Cabral and Riordan (1994) tells us that this is something that would not happen without organizational forgetting: if $\delta = 0$, the equilibrium price of a firm at the bottom of the learning exceeds its marginal cost no matter what the state of its rival firm, i.e., $p^*(m, e_2) > c(m)$ for any $e_2 \in \{1, \dots, M\}$.¹⁵ However, when there is organizational forgetting a firm that reaches the bottom of the learning curve might price below marginal

¹⁵Cabral and Riordan prove their theorem in a model in which there is no outside demand alternative. The logic of their proof of Theorem 4.4 can be readily adapted to the setting in our model where there is an outside alternative. A proof is available on request.

cost. For example, when $\rho = 0.85$, $\delta = 0.08$, $p^*(15, 14) = 4.07$ and $p^*(15, 15) = 3.21$, both of which are less than the firm's marginal cost $c(m)$ of 5.30.

4.3 Wells, trenches, and resultant forces

The most dramatic manifestations of the intensified price competition brought about by organizational forgetting are the wells and trenches in the pricing function. Wells (as seen in the top right-hand panel and the two bottom panels of Figure 2) represent an intense preemption battle between the firms to establish an initial advantage, followed by a quick “surrender” by the follower once one of the firms has established an initial advantage. As this figure (as well as Table 2) shows, even a modest amount of organizational forgetting can give rise to intense races to acquire initial advantage.

Diagonal trenches (as seen in the two bottom panels of Figure 2) correspond to price wars that are triggered when a follower moves from a position of cost disadvantage to one of cost parity. Diagonal trenches are fueled by the prize that accrues to a firm from assuming the role of industry leader and from avoiding the role of industry follower. Put another way, because organizational forgetting makes it more difficult for both firms to move down their learning curve in tandem, diagonal trenches reflect the strong incentive that each firm has to fight hard to move off — and stay off — the diagonal of the state space.

Sideways trenches (as seen in the bottom right-hand panel of Figure 2) are, like diagonal trenches, price wars that are triggered when a follower improves its competitive position. In contrast to a diagonal trench, however, the price war in the sideways trench is not triggered because the follower has moved into a position of cost parity. Instead, it arises in order to keep a severely disadvantaged firm “in its place” as a marginal force in the industry. That is, a sideways trench is designed to prevent a firm with minimal know-how from “breaking through” to a position where it can evolve into a firm that could, at some point in the future, represent a threat to the profitability of the market leader. If diagonal trenches are about fighting against an imminent threat, sideways trenches are about fighting a distantly imagined threat. One might think of them as an equilibrium manifestation of the dictum, “Only the paranoid survive.”

Pricing behavior drives the evolution of the industry. The impact of wells and trenches on equilibrium behaviors can be highlighted through the use of resultant forces diagrams. Given that the industry is in state \mathbf{e} in the current period, it will be in state \mathbf{e}' in the next period, where $e'_k - e_k \in \{-1, 0, 1\}$, $k = 1, 2$. We determine the expected movement of the state by computing the probability weighted average of $\mathbf{e}' - \mathbf{e}$ using the Markov transition probabilities. Reinterpreting $\mathbf{e}' - \mathbf{e}$ as a direction and the associated probability as the force operating in that direction, the expected movement $E(\mathbf{e}' - \mathbf{e}|\mathbf{e})$ becomes the resultant force.¹⁶ Figure 4 shows the resultant forces for the scenarios we have been focusing on: $\rho = 0.85$ and $\delta = 0, 0.03$, and

¹⁶More formally, let $\mu(\mathbf{e}'|\mathbf{e}, p^*(\mathbf{e}))$ denote the Markov transition probability from state $\mathbf{e} = (e_1, e_2)$ to state $\mathbf{e}' = (e'_1, e'_2)$. We first compute the resultant force as

$$E[(\mathbf{e}' - \mathbf{e})|\mathbf{e}] = \sum_{\mathbf{e}' \in \Xi^2} (e'_1 - e_1, e'_2 - e_2)\mu(\mathbf{e}'|\mathbf{e}, p^*(\mathbf{e})).$$

Then, we plot an arrow with the foot at (e_1, e_2) and the head at $(e_1, e_2) + E[(\mathbf{e}' - \mathbf{e})|\mathbf{e}]$.

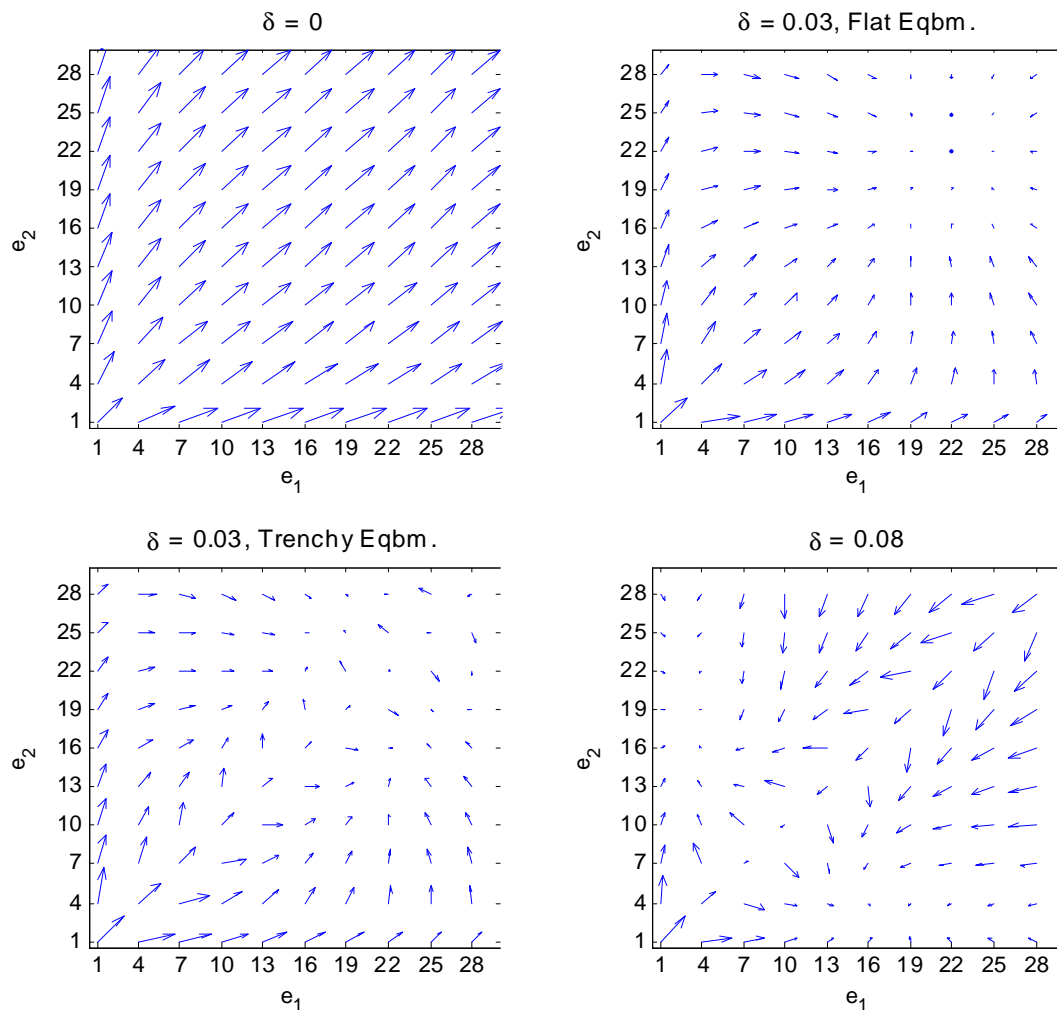


Figure 4: Resultant forces for $\rho = 0.85$ and $\delta = 0, 0.03$ (both equilibria), 0.08 .

0.08 . The arrows in Figure 4 point in the direction in which the industry is expected to evolve, and their length indicates the speed at which this is expected to happen.¹⁷

When there is no forgetting, ($\delta = 0$) the resultant forces indicates that firms tend to acquire know-how in lockstep, with both firms moving steadily to the bottom of the learning curve. The resultant forces in the flat equilibrium with $\delta = 0.03$ shows a similar, but not completely identical, dynamic. In this case, the flow of the resultant forces suggests that the firms would initially be expected to diverge somewhat, a reflection of the existence of the well at $(1,1)$. However, the flatness of the equilibrium beyond this well creates an opportunity for the follower to catch up to the leader, and both firms eventually converge to a symmetric position at the bottom of their learning curves.

¹⁷To simplify the graph, we plot resultant forces only for states that are multiples of 3.

By contrast, in the trenchy equilibrium with $\delta = 0.03$, the resultant forces moves the industry toward points away from the diagonal, a reflection of the diagonal trench in the pricing function. Thus, in contrast to the flat equilibrium, a follower would not be expected to catch up to the leader, suggesting that the industry would tend to evolve to a long-run structure in which firms are asymmetric.¹⁸ The same is true in the extra-trenchy equilibrium with $\delta = 0.08$, but the effect is even more dramatic: the resultant forces suggests that the industry would tend to evolve toward a structure in which one firm is dominant. In this case, we can see the effect of the sideways trench: it makes it unlikely that a follower who operates at the top of the learning curve would be able to “break through” and put itself in a position in which it could challenge the leader.

5 Industry dynamics

Traditional intuition holds that learning-by-doing is likely to lead to market dominance by giving a more experienced firm the ability to profitably underprice its less experienced rival, which enables the leader to widen its cost advantage over time, increasing the scope for underpricing the rival. Cabral and Riordan (1994) formalize this idea with the properties of increasing dominance (ID) and increasing increasing dominance (IID). The MPE exhibits ID if $e_1 > e_2 \Rightarrow p^*(e_1, e_2) < p^*(e_2, e_1)$. The MPE exhibits IID if $p^*(e_1, e_2) - p^*(e_2, e_1)$ decreases in e_1 . If ID holds, the leader charges a lower price than the follower. If IID holds, the gap between the leader’s price and the follower’s price increases as the leader’s cost advantage becomes more pronounced. In a model with no organizational forgetting, Cabral and Riordan prove that ID and IID hold when the discount factor is sufficiently close to 1 and when it is sufficiently close to 0.

Even though the equilibrium pricing function satisfies ID and IID, however, it is not immediately clear that the industry is on an inevitable path toward monopoly or near-monopoly, either in the short run or the long run. The effect of ID and IID on pricing behavior and industry dynamics might be economically insignificant if the difference in the prices of the two firms is small and there is some degree of horizontal differentiation between the firms.¹⁹ In such a scenario, the leader would set a *slightly* lower price than the follower, and this small price gap would grow *a bit* over time as the experience gap widens. But because the difference in prices remains small, with even a modest degree of horizontal differentiation both firms might be able to split the sales more or less equally over time and thus move down the learning curve together. If so, ID and IID would have no discernible impact on the structure of the industry in the short run or the long run.²⁰

¹⁸When we study industry dynamics more explicitly in the next section, we will see that this is in fact the case.

¹⁹Indeed, Cabral and Riordan prove that in the limiting case of $\beta = 1$ both firms charge the same price in every state, i.e., $p^*(e_1, e_2) = p^\dagger(c(m), c(m))$ for all (e_1, e_2) . This suggests that for “reasonable” discount factors the price gap may be small.

²⁰As we will see below, this is exactly what happens when there is no organizational forgetting, or when organizational forgetting is relatively weak.




Figure 5: Timepath of the expected Herfindal index for $\rho = 0.85$ and $\delta = 0, 0.03$ (both equilibria), 0.08 .

5.1 Transient and limiting distributions

Because ID and IID may not be fully informative, we investigate industry dynamics in more detail by using the Markov transition probabilities to compute the transient distribution $\pi^T(\mathbf{e}|(1, 1))$ over states at time T , starting from the initial state $(1, 1)$. To characterize short-run industry dynamics, we compute the expected Herfindahl index implied by the transient distribution at time T ,

$$H^T \equiv \sum_{\tilde{e}_1} \sum_{\tilde{e}_2} \left[\frac{D^*(\tilde{\mathbf{e}})^2 + D^*(\tilde{\mathbf{e}}^{[2]})^2}{[D^*(\tilde{\mathbf{e}}) + D^*(\tilde{\mathbf{e}}^{[2]})]^2} \right] \pi^T(\tilde{\mathbf{e}}|(1, 1)),$$

and to characterize long-run dynamics, we compute the expected Herfindahl index H^∞ associated with the limiting distribution $\pi^\infty(\mathbf{e})$.

Figure 5 plots the time path of H^T for each of our four representative scenarios. To the extent that the firms evolve in a more or less symmetric fashion, the expected industry Herfindahl index should be close to 0.500, and the transient distributions should be unimodal and concentrated along the diagonal of the state space. If, by contrast, the industry structure is asymmetric, the expected industry Herfindahl should exceed 0.500, and the transient distributions should be bimodal with probability weight concentrated closer to the axes.

When $\delta = 0$, the maximal expected Herfindahl index is approximately 0.527 (reached in period 4). Despite the fact that the pricing function is characterized by ID and IID, the industry structure remains highly symmetric as it evolves over time. By $T = 32$, the expected Herfindahl index is essentially 0.500. This case illustrates that that ID and IID are not sufficient for industry dominance, either in the short run or the long run.

In contrast to the case without organizational forgetting, organizational forgetting introduces both short-run and long-run asymmetries in industry structure. Consider, for example, the flat equilibrium for $\delta = 0.03$. As we see from Figure 5, the industry initially becomes asymmetric, attaining an expected Herfindahl index that exceeds 0.650 between periods 2 and 8. Over time, the degree of asymmetry declines, but even after 32 periods, the industry is still mildly asymmetric, with an expected Herfindahl index of 0.514. However, as time passes the asymmetries fade. In the limit, the industry becomes nearly symmetric with $H^\infty = 0.500$. The short-run dynamics for the trenchy equilibrium when $\delta = 0.03$ are similar. However, under this scenario, the asymmetries do not disappear in the long run, with $H^\infty = 0.516$. When $\delta = 0.08$, more exaggerated and persistent asymmetries arise. The industry becomes highly asymmetric right away, attaining an expected Herfindahl index of 0.844 by $T = 8$. However, the degree of asymmetry hardly changes after that, with an expected Herfindahl index of 0.807 and a limiting expected Herfindahl of 0.770. In this case, one firm develops an initial cost advantage that persists over time. Interestingly, even though the likely outcome in this market is dominance by one firm, ID and IID do not hold at all points in the state space (in fact, in this equilibrium, ID is violated in about 14 percent of the states and IID is violated in over 55 percent of the states). Thus, ID and IID are neither necessary nor sufficient for market dominance.

These conclusions are reinforced by direct examination of the transient distributions. Figures 6 and 7 show the transient distributions for $T = 8$ and $T = 32$ for our representative scenarios, and Figure 8 shows the corresponding limiting distribution. (In each panel in Figures 6 - 8, the height of each line is the probability that the industry is in a particular state at that point in time).

The transient distributions for our four scenarios reveal a rich variety of short-run and long-run dynamics. When there is no forgetting, the transient distribution is unimodal and concentrated along the diagonal of the state space. Though temporary asymmetries among firms might arise (e.g., Figure 6 shows that at $T = 8$, there is positive probability that one firm can be in state 9, while the other can still be in state 1), the follower is very likely to catch up, and the asymmetries are unlikely to persist.

By contrast, when there is organizational forgetting, the transient distributions can have very different shapes. In the flat equilibrium for $\rho = 0.85$, $\delta = 0.03$, the transient distributions in the short and intermediate run are bimodal. For example, at $T = 32$, the modes are (15, 7) and (7, 15). In the limit, though, the asymmetries between the firms virtually disappear, with the limiting distribution becoming essentially unimodal. In the trenchy equilibrium for $\rho = 0.85$, $\delta = 0.03$, we get a similar pattern in the long run and the short run (though the two modes are more distinguishable and pronounced), but the asymmetries do not vanish: the limiting distribution has two modes: (25, 18) and (18, 25). In this case, both firms reach

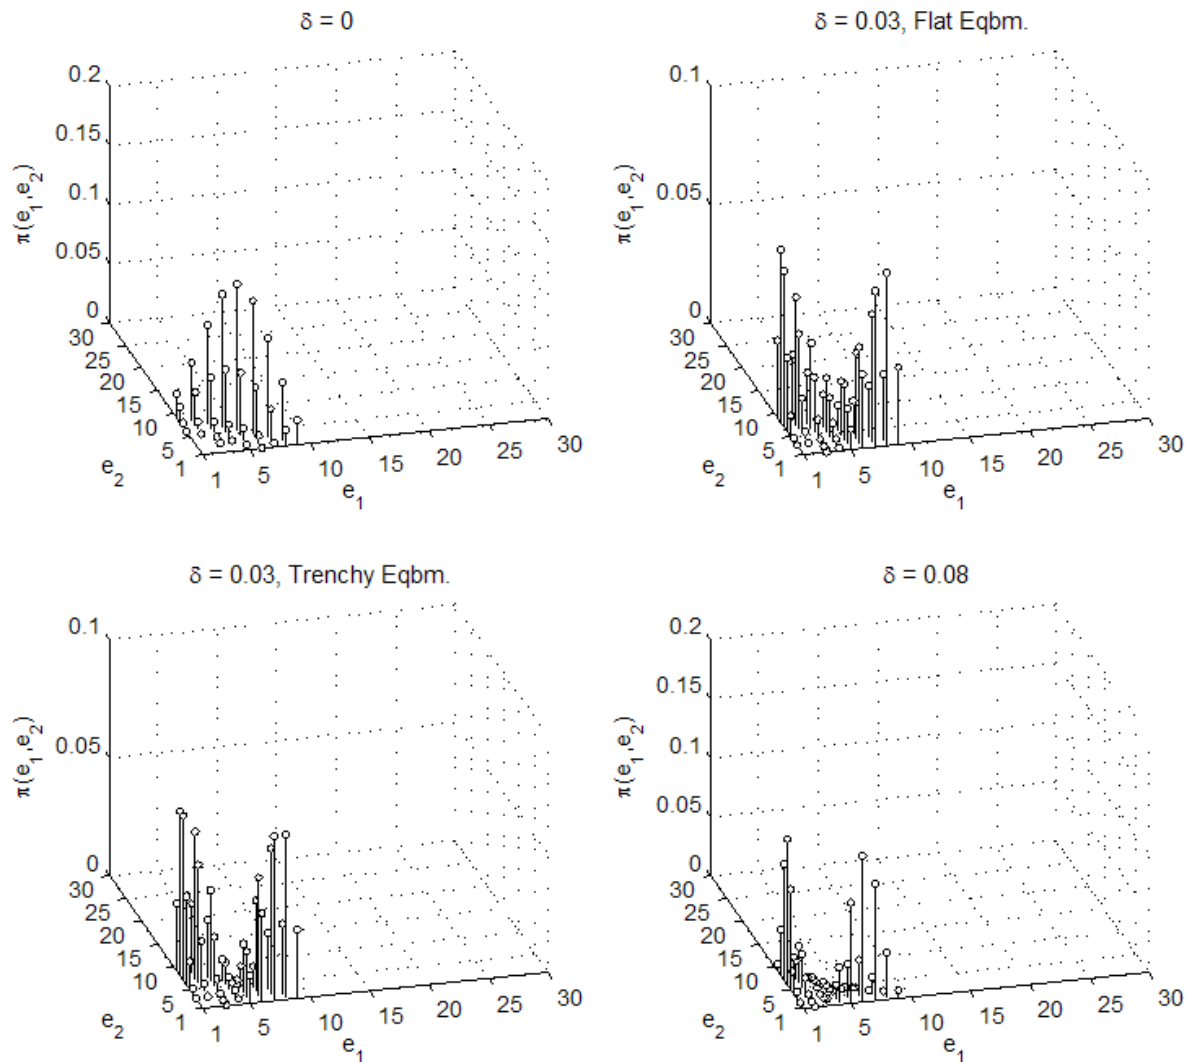


Figure 6: Transient distribution at $T = 8$ (start from $(1, 1)$ at $T = 0$) for $\rho = 0.85$ and $\delta = 0, 0.03$ (both equilibria), 0.08 .

the bottom of the learning curve, but one firm is more vulnerable to forgetting than the other. When $\rho = 0.85, \delta = 0.08$, the transient distributions reveal that the industry essentially evolves along the “edge” of the state space, a reflection of the sideways trenches discussed earlier. As time passes, the modal state of the follower firm remains $e_2 = 1$, while the modal state of the leader moves from $e_1 = 6$ at $T = 8$ to $e_1 = 13$ at $T = 32$ and $e_1 = 19$ as $T \rightarrow \infty$. In this case, the modal outcome is for the leader to make it safely to the bottom of the learning curve, leaving the follower stuck at the top.

Table 3 summarizes the short-run and long-run industry dynamics for all $(\rho, \delta) \in \mathcal{P} \times \mathcal{D}$. It shows the expected Herfindahl Index given the limiting distribution, H_∞^* (repeating the information in Table 1) and the maximum attained expected Herfindahl Index, \widehat{H}^* . Table 4 presents similar information in a slightly different way. It classifies industry structures into categories, depending on the location of the modes of the

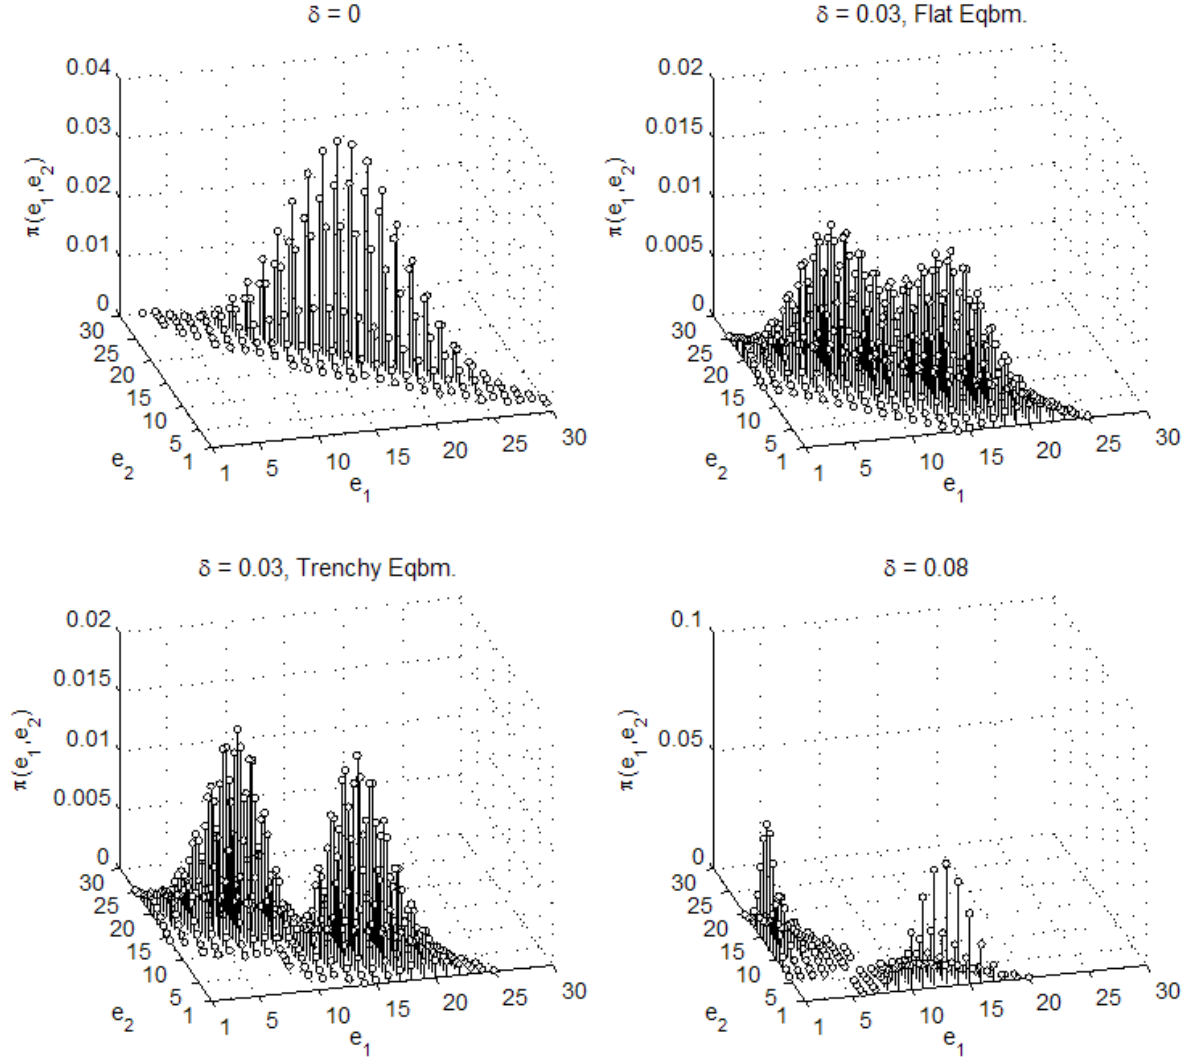


Figure 7: Transient distribution at $T = 32$ (start from $(1, 1)$ at $T = 0$) for $\rho = 0.85$ and $\delta = 0, 0.03$ (both equilibria), 0.08 .

limiting distribution:

- Symmetric duopoly (S): modal state such that $|e_1 - e_2| \leq 1$ and $e_1 > 1, e_2 > 1$.
- Asymmetric duopoly (A): modal state such that $|e_1 - e_2| > 1$ and $e_1 > 1, e_2 > 1$.
- Near-monopoly (M): modal state such that $|e_1 - e_2| > 1$ and $e_1 = 1$ or $e_2 = 1$.
- Stagnant industry (E): modal state is $(1, 1)$.

Tables 3 and 4 suggests three conclusions. First, holding the degree of organizational forgetting fixed, short-run and long-run industry concentration may be non-monotonic in the rate of learning. For example,

ρ	0.95		0.85		0.75		0.65		0.55		0.35		0.15		0.05			
	F	X	F	X	F	X	F	X	F	X	F	X	F	X	F	X		
δ																		
0	.504	.527	.603	—	.671	—	.662	—	.580	—	.550	—	.580	—	.550	—	.544	—
	.500	.500	.500	—	.500	—	.500	—	.500	—	.500	—	.500	—	.500	—	.500	—
0.01	.505	.547	.683	—	.771	—	.753	—	.624	—	.579	—	.624	—	.579	—	.569	—
	.500	.500	.500	—	.500	—	.500	—	.500	—	.500	—	.500	—	.500	—	.500	—
0.02	.507	.600	.798	—	.869	—	.850	—	.679	—	.611	—	.679	—	.611	—	.597	—
	.500	.500	.500	—	.500	—	.500	—	.500	—	.500	—	.500	—	.500	—	.500	—
0.03	.509	.736	.796	—	.873	—	.902	—	.743	.888	.644	.869	.743	.888	.644	.869	.626	.860
	.500	.500	.516	—	.520	—	.521	—	.500	.521	.500	.521	.500	.521	.500	.521	.500	.521
0.04	.511	.825	.823	—	.921	.908	.956	.951	.964	.944	.677	.881	.834	.907	.677	.881	.654	.868
	.502	.502	.513	—	.502	.519	.502	.521	.501	.521	.500	.523	.500	.523	.500	.523	.500	.523
0.05	.513	—	.847	—	—	.936	—	.955	—	—	.963	—	.932	—	.709	.892	.681	.876
	.504	.530	.530	—	.538	—	.580	—	.542	—	.542	—	.527	—	.500	.527	.500	.527
0.06	.515	—	.870	—	—	—	—	.955	—	.963	.964	.956	—	.739	.910	.706	.884	
	.507	.559	.559	—	.882	—	.923	—	.940	.940	.500	.531	—	.500	.530	.500	.530	
0.07	.517	—	—	.881	—	.938	—	.955	—	.963	—	.969	.930	.769	.930	.730	.900	
	.511	.744	.744	—	.884	—	.921	—	.938	.938	.963	.938	.954	.532	.500	.532	.500	.534
0.08	.520	—	—	.882	—	.939	—	.956	—	.963	—	.969	—	.799	—	.754	.876	
	.515	.770	.770	—	.883	—	.920	—	.937	.937	.963	.937	.952	.961	.501	.959	.501	.961
0.09	.522	—	—	.882	—	.939	—	.956	—	.963	—	.969	—	.829	—	.778	—	
	.519	.787	.787	—	.882	—	.918	—	.935	.935	.963	.935	.950	.960	.502	.957	.502	.960
0.1	.524	—	—	.881	—	.939	—	.955	—	.963	—	.969	—	.882	—	.804	—	
	.523	.803	.803	—	.882	—	.916	—	.933	.933	.963	.933	.949	.956	.505	.956	.504	.958
0.3	.517	—	—	.860	—	.935	—	.955	—	.964	—	.971	—	—	—	—	—	
	.517	.860	.860	—	.922	—	.936	—	.942	.942	.964	.942	.945	.945	—	.973	—	.973
0.5	.508	—	—	.597	—	.897	—	.956	—	.966	—	.974	—	—	—	—	—	
	.508	.597	.597	—	.897	—	.956	—	.966	.966	.966	.966	.969	.969	—	.976	—	.976
0.7	.504	—	—	.539	—	.615	—	.761	—	.902	—	.965	—	—	—	—	—	
	.504	.539	.539	—	.615	—	.762	—	.904	.904	.902	.904	.965	.965	—	.975	—	.977
0.9	.501	—	—	.511	—	.532	—	.561	—	.591	—	.668	—	—	—	—	—	
	.501	.511	.511	—	.532	—	.561	—	.591	.591	.591	.591	.672	.672	—	.754	—	.789
																		.803

Table 3: Maximal (top number) and limiting (bottom number) Expected Herfindahl Indices for $(\rho, \delta) \in \mathbf{R} \times \mathbf{D}$.

ρ	0.95		0.85		0.75		0.65		0.55		0.35		0.15		0.05			
δ	<i>F</i>	<i>T</i>	<i>X</i>	<i>F</i>	<i>T</i>	<i>X</i>	<i>F</i>	<i>T</i>	<i>X</i>	<i>F</i>	<i>T</i>	<i>X</i>	<i>F</i>	<i>T</i>	<i>X</i>	<i>F</i>	<i>T</i>	<i>X</i>
0	<i>S</i>	-	-	<i>S</i>	-	-	<i>S</i>	-	-	<i>S</i>	-	-	<i>S</i>	-	-	<i>S</i>	-	-
0.01	<i>S</i>	-	-	<i>S</i>	-	-	<i>S</i>	-	-	<i>S</i>	-	-	<i>S</i>	-	-	<i>S</i>	-	-
0.02	<i>S</i>	-	-	<i>S</i>	-	-	<i>S</i>	-	-	<i>S</i>	-	-	<i>S</i>	-	-	<i>S</i>	-	-
0.03	<i>S</i>	<i>A</i>	-	-	<i>A</i>	-	-	<i>A</i>	-	-	<i>A</i>	-	<i>S</i>	<i>A</i>	-	<i>S</i>	<i>A</i>	-
0.04	<i>S</i>	<i>A</i>	-	<i>SM</i>	<i>AM</i>	-	<i>SM</i>	<i>AM</i>	-	<i>SM</i>	<i>AM</i>	-	<i>S</i>	<i>A</i>	-	<i>S</i>	<i>A</i>	-
0.05	<i>S</i>	-	<i>M</i>	-	<i>A</i>	-	-	<i>A</i>	-	-	<i>A</i>	-	-	<i>A</i>	-	<i>S</i>	<i>A</i>	-
0.06	<i>S</i>	-	<i>A</i>	-	-	<i>M</i>	-	-	<i>M</i>	-	-	<i>M</i>	<i>S</i>	<i>A</i>	-	<i>S</i>	<i>A</i>	-
0.07	<i>S</i>	-	<i>M</i>	-	-	<i>M</i>	-	-	<i>M</i>	-	-	<i>M</i>	-	<i>M</i>	-	<i>S</i>	<i>A</i>	-
0.08	<i>S</i>	-	<i>M</i>	-	-	<i>M</i>	-	-	<i>M</i>	-	-	<i>M</i>	-	<i>M</i>	-	<i>S</i>	<i>A</i>	<i>M</i>
0.09	<i>S</i>	-	<i>M</i>	-	-	<i>M</i>	-	-	<i>M</i>	-	-	<i>M</i>	-	<i>M</i>	-	<i>S</i>	<i>S</i>	<i>M</i>
0.1	<i>S</i>	-	<i>M</i>	-	-	<i>M</i>	-	-	<i>M</i>	-	-	<i>M</i>	-	<i>M</i>	-	<i>S</i>	<i>S</i>	<i>M</i>
0.3	<i>E</i>	-	<i>M</i>	-	-	<i>M</i>	-	-	<i>M</i>	-	-	<i>M</i>	-	<i>M</i>	-	-	-	<i>M</i>
0.5	<i>E</i>	-	<i>E</i>	-	-	<i>M</i>	-	-	<i>M</i>	-	-	<i>M</i>	-	<i>M</i>	-	-	-	<i>M</i>
0.7	<i>E</i>	-	<i>E</i>	-	-	<i>E</i>	-	-	<i>E</i>	-	-	<i>M</i>	-	<i>M</i>	-	-	-	<i>M</i>
0.9	<i>E</i>	-	<i>E</i>	-	-	<i>E</i>	-	-	<i>E</i>	-	-	<i>E</i>	-	<i>E</i>	-	-	-	<i>E</i>

(*S*)ymmetric duopoly, (*A*)symmetric duopoly, Near-(*M*)onopoly, (*E*)stagnant industry

Table 4: Classification of modes in limiting distribution for $(\rho, \delta) \in \mathbf{R} \times \mathbf{D}$.

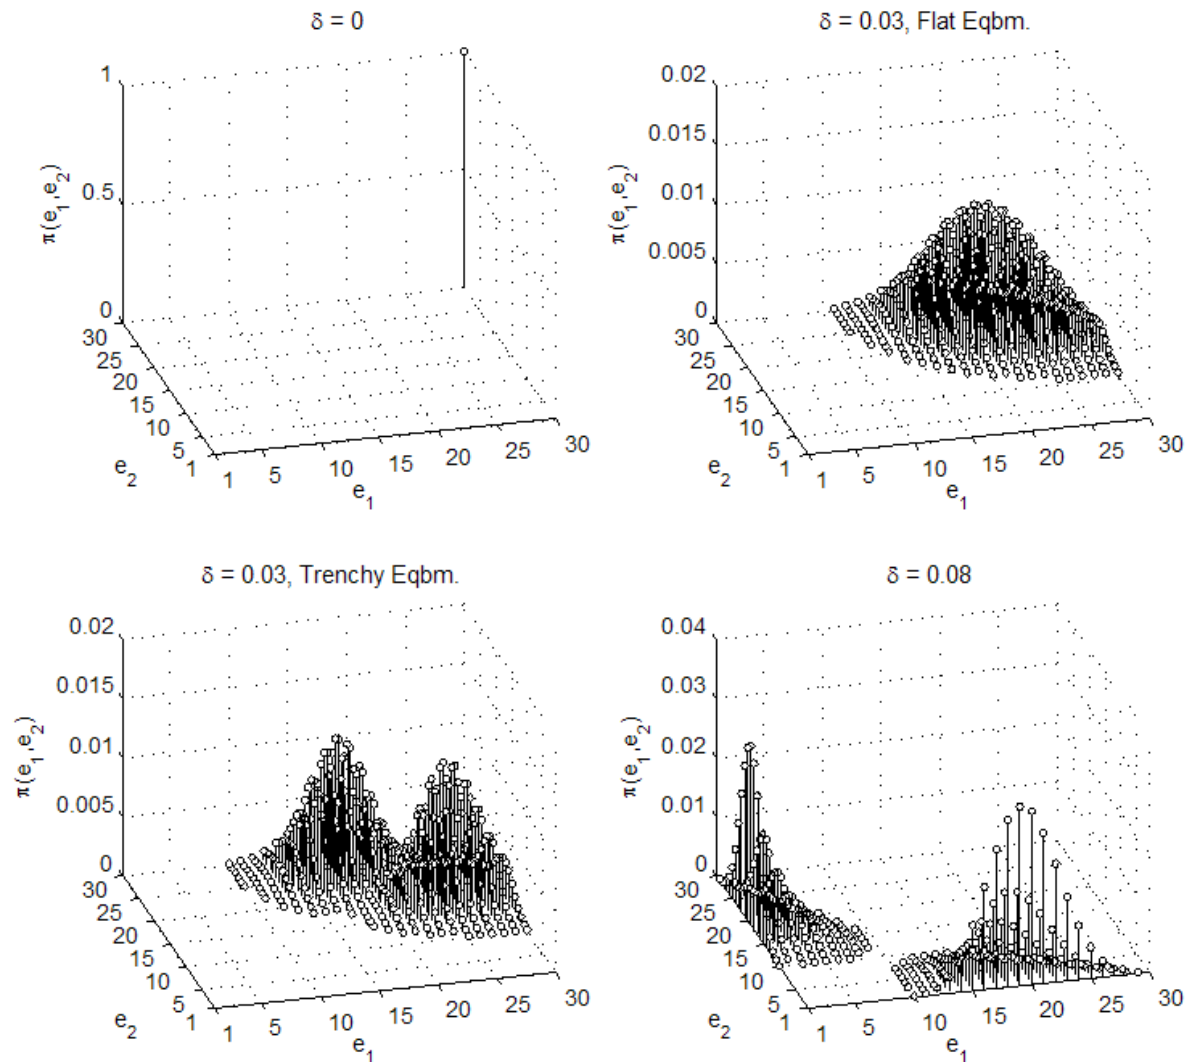


Figure 8: Limiting distribution for $\rho = 0.85$ and $\delta = 0, 0.03$ (both equilibria), 0.08 .

when there is no organizational forgetting, the maximal expected Herfindahl index increases as we move from $\rho = 0.95$ to $\rho = 0.55$, but beyond $\rho = 0.55$ it decreases. We see a similar non-monotonicity in both the maximal and limiting Herfindahl indices when $\delta = 0.05$. These examples tell us that a steeper learning curve does not necessarily give rise to a greater degree of market dominance, either in the short run or the long run. The intuition is this: if the learning curve is sufficiently steep, a firm may be compelled to pursue learning-based cost reductions, no matter how significant its rival's cost advantage. Given this, a cost advantage based on a learning curve may not be sustainable, and the equilibrium market structure may be relatively symmetric.

Second, if organizational forgetting becomes sufficiently severe ($\delta > 0.50$), the firms end up trivially symmetric: they operate at or close to top of the learning curve, with neither firm able to achieve or sustain

an advantage over the other. This reflects the investment-stifling effect of organizational forgetting: when organizational forgetting is severe, it does not make sense to seek learning-based cost reductions.

Third, for small and intermediate values of δ , organizational forgetting tends to make a learning-based cost advantage more sustainable and could even result in a near-monopoly structure in the long run as we saw in the case of $\rho = 0.85, \delta = 0.08$. This conclusion might seem surprising. After all, if an industry leader’s know-how is vulnerable to an increasing rate of organizational forgetting (as in our specification), one might expect that it could never get too far ahead of its rival, and every so often, the leader might drift “back up” its learning curve to a position of parity with its rival. Furthermore, if there is a likelihood that an industry leader’s cost advantage might erode, it might even be worthwhile for a follower to continue to set prices that keep it “close” to the leader and poised to capitalize if the leader’s cost advantage does begin to erode. Thus, one might expect that organizational forgetting would serve as a leveling force that would tend, over time, to result in a symmetric market structure.

But this intuition is incomplete because it ignores the preemption effect of organizational forgetting. As discussed earlier, the preemption effect motivates an industry leader to price aggressively to defend its position against immediate and distant competitive threats. Given this pricing behavior, the industry leader may find that its cost advantage is more secure than it would have been in the absence of organizational forgetting.

To illustrate this last point, we use the transient distribution at each period T to calculate the expected number of periods it will take for a *role reversal*, a point at which the industry moves from an asymmetric state to a symmetric state (or put differently, a point at which the industry moves to the diagonal of the state space). Figure 9 shows expected role reversal time as a function of T for each of our representative scenarios. Because both modes of the limiting distribution are contained in a single recurrent set, role reversals occur from time to time. However, as Figure 9 shows, role reversals become exceedingly unlikely the higher the rate of forgetting. For example, when $\rho = 0.85, \delta = 0$, after 32 periods the expected time for a role reversal is around 30 periods. By contrast, when $\rho = 0.85, \delta = 0.08$, after 32 periods, the expected time for a role reversal is 25,000 periods, essentially forever!

5.2 Organizational forgetting and multiple equilibria

The industry dynamics spawned by organizational forgetting shed light on why multiple equilibria can arise in our model. As we have seen, organizational forgetting can make it extremely difficult for two firms to coexist in a symmetric cost position over time. This sets the stage for a self-fulfilling hypothesis of market dominance:

- Firms *believe* that symmetric coexistence is virtually impossible and that a single dominant firm will emerge, which implies ...
- ... in symmetric states, each firm prices aggressively in order to become the market leader.

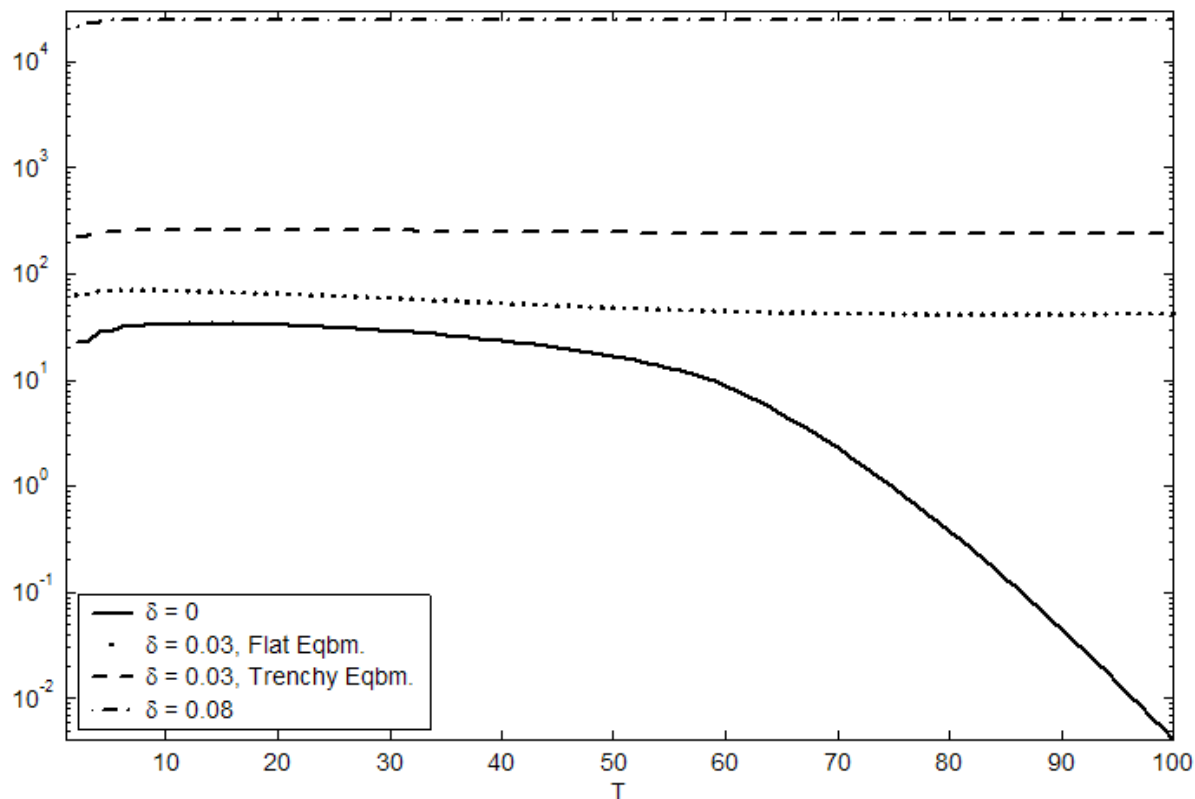


Figure 9: Timepath of expected role reversal time for $\rho = 0.85$ and $\delta = 0, 0.03$ (both equilibria), 0.08 .

- The price wars in symmetric states ensure that symmetric coexistence is, in fact, impossible, thereby leading to the emergence of a dominant firm.

If the firms believed that they *could coexist* in a symmetric cost position over time, they would not fight as aggressively to avoid being symmetric, and they would indeed be able to coexist, i.e., the market would not be winner-take-all. Put another way, organizational forgetting can give the MPE a bootstrap aspect: trenchy equilibria are trenchy because two symmetric firms could not survive “trench warfare” for very long, which in turn makes the prophecy of trench warfare self fulfilling. If firms believed that coexistence is possible, trench warfare would not arise, which indeed makes coexistence possible.

This intuition suggests that multiple equilibria are likely to arise when the rate of organizational forgetting is just large enough to make symmetric coexistence *extremely difficult* in the midst of intense price competition, but *possible* when price competition is sufficiently restrained. This would explain why the multiple equilibria tend to disappear when δ is close to zero or when δ is sufficiently large. When δ is close to zero, symmetric coexistence is possible even when price competition is intense. When δ is sufficiently large, symmetric coexistence is likely to be difficult even when price competition is reasonably mild.

The reason that multiple MPE arise in our model is thus different from the reason for the multiple equilibria in Cabral and Riordan (1994) when they enrich their model to allow for exit. In Cabral and

Riordan, the prospect of driving the other firm out of the market makes it tempting for the industry leader to price aggressively, but the fact that the industry leader prices aggressively, makes it self-fulfilling for the follower to exit. In our model, it is organizational forgetting, not the presence of avoidable fixed costs, which makes it difficult for two equal-sized firms to coexist and gives rise to the possibility of a bootstrap equilibrium in which the firms compete aggressively to achieve a position of industry dominance.²¹

6 Robustness checks

To check the robustness of our conclusions, we investigate how two key parameters — σ (the product differentiation parameter) and $v_0 - c_0$ (attractiveness of the outside alternative) — affect the nature of the MPE. Specifically, for each (ρ, δ) combination in the parameter set $\mathbf{R} \times \mathbf{D}$, we calculated the MPE for each of the following scenarios:

1. Weak horizontal differentiation: $\sigma = 0.10$ and $v_0 - c_0 = 0$ (baseline value).
2. Strong horizontal differentiation: $\sigma = 2$ and $v_0 - c_0 = 0$ (baseline value).
3. (Infinitely) unattractive outside alternative (the case studied by Cabral and Riordan (1994)): $\sigma = 1$ (baseline value) and $v_0 - c_0 = -\infty$.
4. Very attractive outside alternative: $\sigma = 1$ (baseline value) and $v_0 - c_0 = 3$.

Note that increasing the attractiveness of the outside alternative effectively shrinks the market of the inside goods, and when $v_0 - c_0 = 3$, the inside goods have an advantage over the outside alternative only after each firm moves sufficiently far down the learning curve.²²

Tables analogous to Table 3 that report the results of these calculations are available on request. The key conclusions that emerge from this analysis are as follows:

- In all of the above scenarios, we continue to find multiple equilibria.
- The pattern of results that we saw for the case of $\sigma = 1$ and $v_0 - c_0 = 0$ continues to hold: modest amounts of organizational forgetting act as a force that results in less symmetric market structures.
- Holding ρ , δ , and $v_0 - c_0$ fixed, weaker horizontal differentiation tends to result in a more asymmetric market structure. This makes sense. For low values of σ each firm is tempted to engage in aggressive preemptive pricing behavior since undercutting one’s rival by even a tiny amount increases the

²¹As we will see below, adding entry and exit to our model adds this additional “bootstrapping” effect to our model and increases the set of cases in which we have multiple equilibria (including some in which $\delta = 0$).

²²To see this more concretely, note that if $v_0 - c_0 = 3$ and each firm sets a price equal to marginal cost $c(1) = 10$, the probability that a firm makes the next sale would be about 4.5 percent, and the outside good would win 91 percent of the time. By contrast, if each firm reaches the bottom of a 85 percent learning curve and charged the static Nash equilibrium price (which turns out to be 6.84), the probability that a firm makes the next sale would be about 35 percent, and the outside good would win about 30 percent of the time.

probability of winning to close to 1. Even when there is no organizational forgetting, the industry can remain quite asymmetric for many periods. By contrast, when horizontal differentiation is extremely strong, firms can peacefully coexist with each other, and each follows an equilibrium pricing policy that resembles that followed by a monopolist. Unless organizational forgetting is sufficiently strong, the firms move down their learning curves more or less in tandem and attain approximately the same level of cost.

- Elimination of the outside alternative ($v_0 - c_0 = -\infty$) has essentially no impact on our key results. The equilibrium pricing and value functions remain essentially unchanged.
- For sufficiently low values of δ , the increase in $v_0 - c_0$ from 0 to 3 tends to result in more pronounced asymmetries between firms in the short run and possibly even more pronounced asymmetries in the long run. For example, when $v_0 - c_0 = 3$, $\rho = 0.85$, $\delta = 0$, the maximal expected Herfindahl is 0.628, as compared to a maximal expected Herfindahl of 0.527 when $v_0 - c_0 = 0$, $\rho = 0.85$, and $\delta = 0$. However, for sufficiently high values of δ , the increase in $v_0 - c_0$ tends to result in less pronounced short-run and long-run asymmetries. The intuition is this. An increase in the attractiveness of the outside option has offsetting effects that are similar to the investment-stifling and preemption effects due to organizational forgetting. First, by shrinking the market for the inside goods, it becomes more difficult for both firms to move down the learning curve together, making the market more like winner-take-all, thereby intensifying price competition. When δ is low, this tends to deepen the wells and trenches in the equilibrium pricing function, which in turn exaggerates the magnitude and persistence of short-run asymmetries between firms. Offsetting this, however, is that the smaller market size makes investments in the development of a learning-based cost advantage less valuable, reducing firm's willingness to price low when they are the top of the learning curve. This works to reduce the critical value of δ at which the investment-stifling impact of organizational forgetting dominates the preemption effect.

This last point suggests that whether organizational forgetting stifles investment in know-how (leading to an outcome in which both firms are stuck at the top of the learning curve), or intensifies investment in know-how by heating up price competition (leading to an asymmetric oligopoly or a near-monopoly market structure), will depend on the size of the market of the inside goods. In a situation in which movement down the learning curve is a necessary condition for the “inside goods” to attain market penetration, organizational forgetting may operate in the manner suggested by traditional intuition: it acts as a force not unlike a flattening of the learning curve. In a situation in which the inside goods can “win” against the outside alternative even when the inside goods are at the top of their learning curves, organizational forgetting may intensify price competition and induce more asymmetric structures.

7 Entry and exit

So far we have assumed that two firms are present in the market at the outset and never exit. It is straightforward to extend the model to allow for entry and exit. Appendix B presents the formal specification of the model that we use to study entry and exit with multiple firms; here, we will briefly sketch the model.

We assume that there at any given time there is a total number N firms, each of which can either be an incumbent or potential entrant. Thus, If N^* is the number of incumbent firms, $N - N^*$ is the number of potential entrants. If an incumbent firm exits the industry in a given period, it perishes and a potential entrant automatically takes its “slot” and has to decide whether or not to enter the industry in that period. Potential entrants are drawn from a large pool. Hence, if a potential entrant chooses not to enter the industry in a given period, it disappears and another potential entrants takes its slot in the next period. Throughout this section, we focus on the case in which $N = 2$.

To ensure the existence of an equilibrium, we use the approach in Doraszelski and Satterthwaite (2004) in which each potential entrant receives a privately observed “draw” \tilde{S}_n from a distribution of possible set-up costs, and active firms operating in the industry receive a privately observed draw \tilde{X}_n from a distribution of possible salvage values. This specification allows us to characterize entry and exit behavior using an operating probability $\lambda(\mathbf{e}) \in [0, 1]$, where $\mathbf{e} \in \{0, 1, \dots, M\} \times \{0, 1, \dots, M\}$ is the state, and $e_n = 0$ corresponds to the situation in which firm n is a potential entrant. The $\lambda(\mathbf{e})$ of an incumbent firm indicates the probability that the firm remains in the market in a given period. The $\lambda(\mathbf{e})$ of a potential entrant indicates the probability that the entrant comes into the market in the period.

The symmetric MPE consists of a pricing functions $p(\mathbf{e})$ and an operating probability $\lambda(\mathbf{e})$ for the representative firm (where $p = \infty$ if the firm is not active). Though alternative specifications are possible, we focus on the case in which a new entrant comes into the industry at the top of the learning curve.

Table 5 summarizes the equilibrium calculations for all combinations of $(\rho, \delta) \in \{0.85\} \times \mathbb{D}$. In these calculations, we assume that a new entrant pays a set-up cost that is uniformly distributed on the interval $[3, 6]$, while an incumbent firm that exits receives a salvage value that is uniformly distribution on the interval $[0, 3]$.²³

Table 5 shows that the inclusion of entry and exit does not alter the thrust of the results that we described above: organizational forgetting can act as a force that intensifies pricing rivalry and leads to more asymmetric market structures. However, allowing for the possibility that firms can exit adds an additional component to the prize from winning a sale: by winning the next sale, a firm can move the industry into a state in which a rival is more likely to exit, in which case it may be replaced by a firm with an even higher cost, or maybe not replaced at all. As a result, entry and exit and organizational forgetting can interact to increase the depth of trenches, and likelihood of eventual concentration. A comparison of Table 5 with the corresponding column of Table 3 reveals that for any given level of δ , entry and exit makes it more likely

²³Notice that this implies that some portion of entry costs is necessarily sunk. This eliminates the possibility of “speculative entry”: entering the industry in the hope of receiving a draw on the salvage value that exceeds the entry cost.

δ	F	T	X
0	.527 .500	—	—
0.01	.547 .500	—	—
0.02	.600 .500	—	—
0.03	.765 .500	—	.945 1.000
0.04	.857 .502	—	.999 1.000
0.05	—	—	1.000 1.000
0.06	—	—	1.000 1.000
0.07	—	—	1.000 1.000
0.08	—	—	1.000 1.000
0.09	—	—	1.000 1.000
0.1	—	—	1.000 1.000
0.3	—	—	.997 .997
0.5	—	—	.659 .597
0.7	—	—	.690 .539
0.9	—	—	.728 .511

Table 5: Maximal (top number) and limiting (bottom number) Expected Herfindahls for $(\rho, \delta) \in \{0.85\} \times \mathbb{D}$, with classification of equilibria.

that we will see a trenchy or extra-trenchy equilibrium, and the long-run market structure is likely to be more concentrated.

7.1 Predatory pricing

In the model with entry, exit, and organizational forgetting, equilibrium pricing behavior can give rise to behavior that seems to resemble traditional notions of predatory pricing. In the trenchy and extra-trenchy equilibria that arise for intermediate values of δ , competition between the two firms plays out something like this: if both firms enter the market, they fight an intense price war as long as both remain at the top of the learning curve. If one firm gains a cost advantage by moving one or more steps down the learning curve, both firms may continue to charge prices below their marginal costs, but the leader now under-prices the follower. If the leader's cost advantage become sufficiently large, the follower eventually exits, after which

the leader increases its price substantially, up to a level roughly equal to the static monopoly price for the marginal cost at the bottom of the learning curve.

In this situation, can the low-price firm plausibly be accused of engaging in predatory pricing? In other words, does predatory pricing arise in the equilibrium of our model? To explore this question, we adapt two “economic definitions” of predatory pricing that have been proposed in the literature: one by Ordover and Willg (1981); the other by Cabral and Riordan (1997).

Ordover and Willg (1981) define a predatory action as “a response to a rival that sacrifices part of the profit that could be earned under competitive circumstances, were the rival to remain viable, in order to induce exit and gain consequent monopoly profits.” An operationalization of that test in our model would be to determine whether the low-price firm’s price is lower than it would have been in a counterfactual model in which firms are assumed to remain viable, i.e., to never exit.²⁴ Formally:

Definition 2 (*Ordover-Willig Test of Predatory Pricing*) Let $p^*(\mathbf{e})$ be the pricing function in a model with entry and exit, and let $\min\{p^*(\mathbf{e}), p^*(\mathbf{e}^{[2]})\}$ be the price charged by the low-price firm in states $\chi \equiv \{\tilde{\mathbf{e}} | \tilde{e}_1 > \tilde{e}_2\}$ when the firms are asymmetric.²⁵ Further, let $p^{*0}(\mathbf{e})$ be the equilibrium price that the low-price firm would have charged in the counterfactual model without exit. The low-price firm’s price is predatory in state $\mathbf{e} \in \chi$ if $\min\{p^*(\mathbf{e}), p^*(\mathbf{e}^{[2]})\} < p^{*0}(\mathbf{e})$, i.e., the low-price firm’s price is lower than it would have been in the absence of exit.^{26,27}

A potential limitation of this definition of predatory pricing is that the “never exit” counterfactual is quite strong. A firm in our model may exit because it happened to get a particularly high realization of the salvage value \tilde{X}_n , and the fact that it exits under these circumstances does not mean that its rival intended to force it out of the market through predatory behavior. A second (even more subtle) limitation is that in our model with exit, a firm may charge a lower price than it would have charged in a model without exit, not because it intends to hasten the exit of its rival in the next period (the classic notion of a predatory action), but rather because it seeks to improve its own future competitive position in the hope that it may eventually come to pass that its rival exits at some distant point in the future.

In light of these limitations, we study an alternative definition of predatory behavior offered by Cabral and Riordan (1997): “an action is predatory if (1) a different action would increase the likelihood that rivals

²⁴It might be the case that the high-price firm’s price is also lower than it would have been in the absence of exit. However, this price would not be predatory since presumably the high-price firm could argue that its price was intended to meet the competition.

²⁵There is no loss of generality in focusing on states in which $e_1 > e_2$. Given the symmetry of the MPE, states in which $e_2 > e_1$ are “mirror images” of states in which $e_1 > e_2$.

²⁶In symmetric states ($e_1 = e_2 = e$) it might also be the case that $p^*(e, e) < p^{*0}(e, e)$. However, in these states the firms charge the same price, so there is no distinction between the predator and the prey. It seems unlikely that antitrust policy would penalize firms in this situation, and so these cases are not classified as predatory pricing.

²⁷A complication in applying this definition arises if we have multiple equilibria. We deal with this by following the convention of comparing “like” equilibria. For example, if a model with entry and exit there is a flat and a trenchy equilibria, and there is a corresponding flat and trenchy equilibrium in the model without entry and exit, we compare the flat equilibrium to the flat equilibrium, and the trenchy equilibrium to the trenchy equilibrium. If the model with entry and exit has two or more equilibria, but there is a unique equilibrium in the model without entry and exit, we compare all equilibria to the unique equilibrium without entry and exit. The opposite (more problematic) case where there are more equilibria without entry and exit than there are with entry and exit did not arise in our calculations.

would remain viable; (2) the different action would be more profitable under the counterfactual hypothesis that the rival’s viability were unaffected.” Cabral and Riordan formalize their definition using a two-period model in which one firm can take an action in the first period that increases the likelihood that a rival firm will exit at the beginning of the second period. In such a model, an action would be predatory if a *lower* action would be *more profitable* under the counterfactual hypothesis that the rival firm’s probability of exit is fixed at its equilibrium level. This test differs from the Ordover-Willig test in that it allows for the possibility that there could be a positive probability that the rival exits in equilibrium.²⁸

To operationalize the Cabral Riordan test in our model, we use the following counterfactual. Consider a model in which there can be entry and exit, but imagine that a firm acts “as if” the industry *cannot* transition to a state in which its rival exits next period. This forces the firm to “ignore” the “prize” that it gets if the industry transitions to a state in which follower exits in the next period, which, in turn, implies that the firm’s current pricing decision will not internalize this component of the prize. Unlike the Ordover-Willig “no exit” counterfactual, however, the rival in this counterfactual can still exit if it is optimal to do so. Note that this counterfactual focuses on exit inducement in the *next period*. A firm in this counterfactual *is permitted* to internalize the incremental benefit of transitioning to a state from which, at some distant point in the future, the industry may reach *other states* in which the rival may exit. In this sense, this counterfactual gives the firm more latitude to price aggressively than in the “no exit” counterfactual.

We formalize the Cabral-Riordan standard for predatory pricing as follows:

Definition 3 (*Cabral-Riordan Test of Predatory Pricing*) Let $p^*(\mathbf{e})$ be the pricing function in a model with entry an exit, and let $\min\{p^*(\mathbf{e}), p^*(e^{[2]})\}$ be the price charged by the low-price firm in states $\chi \equiv \{\tilde{\mathbf{e}}_1 > \tilde{\mathbf{e}}_2\}$ when the firms are asymmetric. Further, let $\hat{p}^*(\mathbf{e})$ be the equilibrium price that the low-price firm would have charged if it had ignored the impact of its pricing decision on its rival’s next-period exit behavior.²⁹ The low-price firm’s price is predatory in state $\mathbf{e} \in \chi$ if $\min\{p^*(\mathbf{e}), p^*(e^{[2]})\} < \hat{p}^*(\mathbf{e})$, i.e., the low-price firm’s price is lower than it would have been had it not ignored the impact of its pricing decision on the rival’s next-period exit behavior.

Table 6 and Table 7 summarize the extent of predatory pricing according to each of the two standards when $\rho = 0.85$.³⁰ Each table reports three numbers. The first is the percentage of asymmetric states in which the low-price firm engages in above-cost predatory pricing. The second is the percentage of asymmetric states in which the low-price firm engages in below-cost predatory pricing (and so the sum of these two percentages is the overall incidence of predatory pricing). The third is the percentage of asymmetric states in which the low-price charges a below-cost price that is non-predatory.

For the set of parameter values in Tables 6 and 7, the two standards yield different results. As we would

²⁸In the two-period model studied by Cabral and Riordan, this test can be shown to differ from the Ordover and Willig test.

²⁹In Appendix 2, we discuss how we alter the model with entry and exit to compute this counterfactual equilibrium price $\hat{p}^*(\mathbf{e})$.

³⁰Throughout this analysis, we assume that as above, a new entrant pays a set-up cost that is uniformly distributed on the interval $[3, 6]$, and an incumbent firm that exits receives a salvage value that is uniformly distribution on the interval $[0, 3]$.

	Predatory pricing:	Predatory pricing:	Non-predatory pricing:
δ	above cost	below cost	below cost
0	0%	0%	0.7%
0.01	0%	0%	0.7%
0.02	0%	0%	0.7%
0.03 (<i>F</i>)	19.8%	0%	1.1%
0.03 (<i>T</i>)	11.5%	7.8%	1.6%
0.04 (<i>F</i>)	0%	0.2%	2.3%
0.04 (<i>T</i>)	8.5%	14.3%	1.8%
0.05	23.7%	13.1%	1.4%
0.06	49.7%	8.5%	0.5%
0.07	50.6%	7.6%	0%
0.08	60.5%	9.2%	0%
0.09	63.7%	10.1%	0%
0.10	62.8%	10.6%	0.2%
0.30	38.9%	9.2%	0.7%
0.50	33.1%	7.1%	2.3%
0.70	31.5%	6.4%	3.0%
0.90	29.4%	6.4%	3.0%

Table 6: Incidence of predatory pricing under the Orderover-Willig test when $\rho = 0.85$

	Predatory pricing:	Predatory pricing:	Non-predatory pricing:
δ	above cost	below cost	below cost
0	0%	0%	0.7%
0.01	0%	0%	0.7%
0.02	0%	0%	0.7%
0.03 (<i>F</i>)	15.6%	0%	1.1%
0.03 (<i>T</i>)	11.5%	7.8%	1.6%
0.04 (<i>F</i>)	0%	0.2%	2.3%
0.04 (<i>T</i>)	3.4%	13.6%	1.8%
0.05	18.2%	10.6%	1.4%
0.06	40.5%	7.1%	0.5%
0.07	39.1%	7.6%	0%
0.08	48.5%	9.2%	0%
0.09	52.6%	10.1%	0%
0.10	53.8%	10.6%	0.2%
0.30	38.2%	9.2%	0.7%
0.50	32.9%	7.1%	2.3%
0.70	31.3%	6.4%	3.0%
0.90	29.2%	6.4%	3.0%

Table 7: Incidence of predatory pricing under the Cabral-Riordan test when $\rho = 0.85$

expect, the Cabral-Riordan standard is more restrictive than the Ordover-Willig standard. For example, when $\rho = 0.85, \delta = 0.05$, in almost 37 percent of the asymmetric states, the Ordover-Willig test identifies predatory pricing, while the Cabral-Riordan test identifies predatory pricing in just about 29 percent of the asymmetric states.

That said, there are several broad conclusions that are implied by either standard. First, predatory pricing can indeed arise in equilibrium, but the predatory price need not be below marginal cost. For example, when $\delta = 0.08$, under the Cabral-Riordan standard, in more than 48 percent of the asymmetric states we have above-cost predatory pricing, while in about 9 percent of these states we have below-cost predatory pricing.

Second, it is possible to have below-cost pricing that is not predatory. For example, when $\delta = 0$, we have no instances of predatory pricing under either standard, but in 0.70 percent of the asymmetric states we have below-cost pricing. Clearly, then, below-cost pricing is neither a necessary nor a sufficient condition of predatory pricing. Still, the tables suggests that below-cost prices that are non-predatory are rare. Thus, for these parameter values, the incidence of below-cost pricing *usually* is sufficient to indicate predatory behavior under either standard.

Finally, under either standard, the incidence of predatory behavior tends to increase in δ for low levels of organizational forgetting and decreases in δ when δ becomes sufficiently large. When $\rho = 0.85$, the incidence of predatory pricing under the Ordover-Willig standard is greatest when $\delta = 0.09$ (with the low-price firm engaging in predatory behavior in 73.8 percent of the asymmetric states). Under the Cabral-Riordan standard, the incidence of predatory behavior is maximized when $\delta = 0.10$ (with the low-price firm engaging in predatory behavior in over 64 percent of the asymmetric states). By contrast, when $\delta = 0, 0.01$, and 0.02 , there is no predatory behavior under either standard, while at the other extreme, when $\delta = 0.90$, there is predatory behavior in about 35 percent of the asymmetric states under either standard. There is a tendency for predatory behavior to be greatest when, in the absence of exit, organizational forgetting would spawn trenchy or extra-trenchy equilibria in which the probability of re-entry is zero (or close to it) once the follower exits. Under these conditions, it is particularly tempting for the industry leader to price aggressively in order to push the follower firm out of the industry.

7.2 Limit pricing

Limit pricing refers to a situation in which a firm threatened with entry charges a lower price than it would have had it not been threatened with entry. When there is learning-by-doing, limit pricing is potentially rational because by pricing aggressively in the current period, the firm may accelerate the rate at which it moves down its learning curve, thereby making it less attractive for another firm to enter the market later on.

Just as we can use our model to test for predatory pricing, we can also use it to test for the presence of equilibrium limit pricing. In our two-firm model with entry and exit, limit pricing arises when the equilibrium

price $p^*(e_1, 0)$ charged by the firm when it is the sole incumbent is strictly less than the price $p^M(e_1)$ that it would have charged had it been a monopolist that faces a learning curve but is unthreatened by the possibility of future entry. In Table 8 below, we illustrate the incidence of limit pricing for $\rho = 0.85$. As in the tables pertaining to predatory pricing, Table 8 reports three numbers. The first is the percentage of single-firm states (i.e., $(e_1, 0)$) in which the incumbent firm engages in above-cost limit pricing (i.e., $c(e_1) \leq p^*(e_1, 0) < p^M(e_1)$). The second is the percentage of single-firm states in which the incumbent engages in below-cost limit pricing (and so the sum of these two percentages is the overall incidence of limit pricing). The third is the percentage of single-firm states in which the incumbent does not engage in limit pricing (i.e., $p^*(e_1, 0) \geq p^M(e_1)$, but $p^*(e_1, 0) < c(e_1)$).

	Limit pricing:	Limit pricing:	Non-predatory pricing:
δ	above cost	below cost	below cost
0	0%	0%	3.3%
0.01	16.7%	3.3%	0%
0.02	13.3%	3.3%	0%
0.03 (<i>F</i>)	20.0%	3.3%	0%
0.03 (<i>T</i>)	3.3%	3.3%	0%
0.04 (<i>F</i>)	43.3%	3.3%	0%
0.04 (<i>T</i>)	3.3%	3.3%	0%
0.05	6.7%	3.3%	0%
0.06	6.7%	3.3%	0%
0.07	6.7%	3.3%	0%
0.08	6.7%	3.3%	0%
0.09	6.7%	3.3%	0%
0.10	6.7%	3.3%	0%
0.30	26.7%	3.3%	0%
0.50	46.7%	0%	0%
0.70	36.7%	0%	0%
0.90	33.3%	0%	0%

Table 8: Incidence of limit pricing when $\rho = 0.85$

Table 8 indicates that limit pricing can arise in equilibrium, and its incidence is greatest for low levels of δ and for high levels of δ . The intuition is this. Recall from Table 5, that for intermediate values of δ (δ between 0.03 and 0.10), we get trenchy equilibria in which the short-run and long-run structure of the industry is monopolized. If only one firm initially enters the market, then for these parameter values, the equilibrium probability of entry $\lambda^*(0, e_1)$ tends to be positive only when $e_1 = 1$; once the leader moves to states $e_1 \geq 2$, the probability of further entry drops to zero, putting the incumbent in a position where it has a (nearly) secure monopoly.³¹ This corresponds to what Bain (1956) called *blockaded entry*.

Now, for sufficiently large values of δ , even though there may be a zero probability of entry in state $(e_1, 0)$, $e_1 > 1$, there is a non-trivial probability that because of organizational forgetting, an incumbent with know-how $e_1 > 1$ will eventually drift back to $e_1 = 1$, where it will once again be threatened by entry.

³¹We say “nearly” secure because there is, of course, the possibility that the firm could “forget” and eventually slide backward to state $(1, 0)$ where once again it would be threatened by entry.

Hence, for sufficiently large δ , an incumbent will charge a lower price than a monopolist would, *even if the incumbent is not currently in a state in which there is a positive probability of entry*. In other words, the incumbent may engage in limit pricing in states where entry is blockaded in order to reduce the chance of moving back to states where entry is not blockaded.

For the flat equilibria that prevail when δ sufficiently low, the situation is very different. In these cases, there is a positive probability of entry that gets smaller as the leader's experience grows, i.e., $\lambda^*(0, e_1)$ decreases in e_1 . Given this, the incumbent can benefit by setting a price below the price that a monopolist would set in order to accelerate its movement down the learning curve and make entry less likely. This is a classic case of limit pricing.

8 Summary and conclusions

This paper applies the Markov-perfect equilibrium framework presented in Ericson and Pakes (1995) to explore how the presence of a learning curve and organizational forgetting can affect price competition and the evolution of industry structure. We compute symmetric MPE and show that it is possible to obtain multiple MPE for a given set of parameter values. We also saw that the MPE can involve price wars that serve a variety of different purposes: to gain an initial advantage (wells), to ward off immediate threats to one's competitive position (diagonal trenches), to ward off distant threats to one's competitive position (sideway's trenches), or to keep a firm in a position in which it is eventually likely to exit (sideways trenches in a model with entry and exit). Surprisingly, we find that organizational forgetting can often act as a force that intensifies pricing rivalry and makes industry structure more asymmetric. The basic intuition is that organizational forgetting can turn a market in which two or more firms could have peacefully coexisted at the bottom of the learning curve into a winner-take-all or winner-take-most market, which in turn creates an incentive for each firm to fight hard to move ahead and stay ahead of its rival in the race to accumulate know-how. This intuition continues to apply when the industry consists of more than two firms. In that case, price wars arise in equilibrium as part of the shakeout that propels the industry toward its eventual equilibrium structure.

The model in this paper can be extended to analyze a variety of applied and policy issues. For example, the evolution of industry structure is likely to be affected in important ways by the extent of know-how spillovers across firms. This suggests that it might be useful to explore the welfare implications of policies aimed at making it easier for followers or new entrants to benefit from the know-how that industry leaders have generated. The model can also be used to study the impact of policies aimed at subsidizing new entrants or lagging incumbents.

The model can also be extended to incorporate additional economic forces that might play an important role in the evolution of industries in which learning and forgetting are an important part of the business landscape. For example, in the current paper, we have not explicitly modeled capacity constraints or order

backlogs, nor have we explored the role of uncertainty about the eventual size of the markets. These are factors that might be expected to diminish firms' incentives to engage in preemptive pricing behavior, and we would therefore expect that they could have an impact on the equilibrium structure of the industry. In future work, we intend to study these complications.

A Correspondence between capital stock model of learning and forgetting and the model used in this paper

The capital stock model of learning and forgetting is given by the deterministic relationship

$$e_{t+1} = (1 - \xi)e_t + y_t.$$

where ξ is the rate of forgetting per period, and y_t in the flow of orders in a particular period t . Notice that if $y_t = 0$, then if $\xi > 0$, $e_{t+1} < e_t$. Further, note that the higher is e_t increases, y_t must also increase in order to offset the proportional depreciation of know how. If the firm has a steady flow y^* of orders over time, the steady state level of know-how is given by:

$$e^* = \frac{y^*}{\xi}.$$

To illustrate, consider Benkard's (2000) analysis of learning and forgetting in the wide-body airframe market. There were 250 L-1011 aircraft produced over a 14 year period. Assuming a smooth flow of orders, this implies that $q \approx 1.5$ aircraft per month. Benkard found that $\xi \approx 0.04$ per month. This implies that the steady state level of know-how was equal to $\frac{1.5}{0.04} = 37.50$ units.

Now consider our stochastic model of learning and forgetting:

$$\tilde{e}_{\tau+1} = e_{\tau} + \tilde{q}_{\tau} - \tilde{f}_{\tau}.$$

(Since the length of a period in our model is, by definition, the period of time in which there is, at most, one sale, our period length need not correspond to the length of a period in the capital stock model; hence we use of a different time subscript.) Taking expectations gives us:

$$E(\tilde{e}_{\tau+1}|e_{\tau}) = e_{\tau} + \gamma - \delta(e_{\tau})$$

where $\gamma = \Pr(\tilde{q}_{\tau} = 1)$ is the probability that the firm makes a sale in a period, and $\delta(e_{\tau}) = 1 - (1 - \delta)^{e_{\tau}}$. In our model, the steady state level of experience e^* is given by:

$$E(\tilde{e}_{\tau+1}|e^*) = e^*,$$

which implies that:

$$e^* = \frac{\ln(1 - \gamma)}{\ln(1 - \delta)}.$$

We now ask: What is the value $\bar{\delta}$ of the forgetting parameter so that our specification generates a steady state level of experience equal to that implied by an particular capital stock model? The answer is

$$\frac{y^*}{\xi} = \frac{\ln(1 - \gamma)}{\ln(1 - \delta)},$$

or

$$\bar{\delta} = 1 - (1 - \gamma)^{\frac{\xi}{y^*}}$$

To illustrate, consider the data in Benkard's study. If a firm's rate of sales is 1.5 units per month and it faces a monthly depreciation rate of 4 percent, $\frac{\xi}{y^*} = \frac{0.04}{1.5} = 0.02667$. This implies that $\bar{\delta}$ falls in the range between 0 and 0.077 as γ ranges between 0 and 0.95, with $\bar{\delta} \approx 0.018$ when $\gamma = 0.5$. As another example, consider a market in which a firm sells 8 units per month and the monthly depreciation rate is 25 percent (roughly the situation facing

builders of Liberty Ships in Argote and Epple (1990)). This implies a steady-state stock of experience of 32 units. To achieve a comparable steady-state experience in our model in which a firm has a probability γ of making a sale in a given period would require that δ range between 0 and 0.089 as γ ranges between 0 and 0.95, with $\delta \approx 0.021$ when $\gamma = 0.5$.

B General model with entry and exit

In this appendix, we sketch the N -firm version of our model with entry and exit.

B.1 Order of moves

1. Each of the N^* incumbent firms learns its own scrap value and makes an exit decision. Each of the $N - N^*$ potential entrant learns its own set-up cost and makes an entry decision. Entry and exit decisions are made simultaneously. In this process, the industry transits from state \mathbf{e} to state \mathbf{e}' . Specifically, incumbent firm n transits from state $e_n \neq 0$ to state $e'_n = 0$ upon exiting and potential entrant n transits from state $e_n = 0$ to state $e'_n = e^0 \neq 0$ upon entering the industry.
2. Price competition takes place among active firms, where firm n is active if and only if $e'_n \neq 0$. Know-how accumulates and organizational forgetting occurs. In this process, the industry transits from state \mathbf{e}' to state \mathbf{e}'' .

Before making their entry and exit decisions, all firms observe state e , and all firms observe state e' prior to making their pricing decisions.

B.2 Entry and exit

Before price competition takes place, incumbent firms can choose to exit the industry, while potential entrants can choose to enter it. If an incumbent firm exits the industry, it receives a scrap value and perishes. We assume that at the beginning of each period, incumbent firm n draws a random scrap value \tilde{X}_n from a distribution $G_X(\cdot)$ with $E(\tilde{X}_n) = \bar{X}$. The distribution $G_X(\cdot)$ is uniform over the support $[\bar{X} - a, \bar{X} + a]$, where $0 < a \leq \bar{X}$ is a parameter. Thus,

$$G_X(X_n) = \begin{cases} 0 & \text{if } X_n < \bar{X} - a, \\ \frac{1}{2} + \frac{X_n - \bar{X}}{2a} & \text{if } X_n \in [\bar{X} - a, \bar{X} + a], \\ 1 & \text{if } X_n > \bar{X} + a. \end{cases}$$

Scrap values are iid across firms and periods, and firm n 's realized scrap value is observed only by itself but not by its rivals. Let $\tau_n(\mathbf{e}, X_n) = 1$ denote the decision of incumbent firm n to remain in the industry in state \mathbf{e} when it has drawn a scrap value X_n , while $\tau_n(\mathbf{e}, X_n) = 0$ denotes the decision to exit.

Simultaneous with the exit decisions of incumbent firms, potential entrants make entry decisions. If a potential entrant decides not to enter, it receives nothing and perishes. If potential entrant n enters, it incurs a set-up cost. At the beginning of each period, potential entrant n draws a random set-up cost \tilde{S}_n from a distribution $G_S(\cdot)$ with $E(\tilde{S}_n) = \bar{S}$. The distribution $G_S(\cdot)$ is uniform over the support $[\bar{S} - b, \bar{S} + b]$, where $0 < b \leq \bar{S}$ is a parameter. Thus,

$$G_S(S_n) = \begin{cases} 0 & \text{if } S_n < \bar{S} - b, \\ \frac{1}{2} + \frac{S_n - \bar{S}}{2b} & \text{if } S_n \in [\bar{S} - b, \bar{S} + b], \\ 1 & \text{if } S_n > \bar{S} + b. \end{cases}$$

The set-up costs are independently and identically distributed across firms and periods, and its realized set-up cost is private to a firm. Let $\tau_n(\mathbf{e}, S_n) = 1$ denote the decision of potential entrant n to enter the industry in state \mathbf{e} when it has drawn a set-up cost S_n , while $\tau_n(\mathbf{e}, S_n) = 0$ denotes the decision to stay out.

Combining the firms' entry and exit decisions, let $\lambda_n(\mathbf{e})$ denote the probability that firm n operates in the industry in state \mathbf{e} , and $\lambda(\mathbf{e}) = (\lambda_1(\mathbf{e}), \dots, \lambda_N(\mathbf{e}))$ denote the vector of operating probabilities for all firms in the industry. If firm n is an incumbent (i.e., $e_n \neq 0$), then $\lambda_n(\mathbf{e}) = \int \tau_n(\mathbf{e}, X_n) dG_X(X_n)$. If firm n is an entrant (i.e., $e_n = 0$), then $\lambda_n(\mathbf{e}) = \int \tau_n(\mathbf{e}, S_n) dG_S(S_n)$.

B.3 Bellman equations

To develop the Bellman equations, we first consider price setting by active firms. We then consider the exit decisions of incumbent firms and the entry decisions of potential entrants. Throughout this section, we will use the following notation:

- $V_n(\mathbf{e})$ denotes the expected net present value of future cash flows to firm n in state \mathbf{e} before entry and exit decisions have been made.
- $U_n(\mathbf{e}')$ denotes the expected net present value of future cash flows to active firm n in state \mathbf{e}' after entry and exit decisions have been made.

B.3.1 Pricing decisions by active firms

Consider an industry that, via a process of entry and exit, has transitioned from state \mathbf{e} to state \mathbf{e}' . the net present value of future cash flows to active firm n is given by:

$$U_n(\mathbf{e}') = \max_{p_n} \left\{ (p_n - c(e'_n))D_n(p_n, \mathbf{p}_{-n}(\mathbf{e}')) + \beta \sum_{k=0}^N D_k(p_n, \mathbf{p}_{-n}(\mathbf{e}')) \bar{V}_{nk}(\mathbf{e}') \right\}, \quad (\text{B.1})$$

where $p_{-n}(\mathbf{e}')$ denotes the prices charged by the other firms and

$$\begin{aligned} \bar{V}_{nk}(\mathbf{e}') &\equiv E[V_n(\mathbf{e}'') | \mathbf{e}', \text{firm } k \text{ wins the sale}] \\ &= \sum_{e'_1=e_1-1}^{e'_1} \cdots \sum_{e''_{k-1}=e'_{k-1}-1}^{e'_{k-1}} \sum_{e''_k=e'_k}^{e'_k+1} \sum_{e''_{k+1}=e'_{k+1}-1}^{e'_{k+1}} \cdots \sum_{e''_N=e'_N-1}^{e'_N} \\ &\quad \prod_{i \neq k} \Pr(e''_i | e'_i, l) \Pr(e''_k | e'_k, w) V_n(\mathbf{e}''), \\ \bar{V}_{n0}(\mathbf{e}') &\equiv E[V_n(\mathbf{e}'') | \mathbf{e}', \text{no firm wins the sale}] \\ &= \sum_{e'_1=e_1-1}^{e'_1} \cdots \sum_{e''_N=e'_N-1}^{e'_N} \prod_{i=1}^N \Pr(e''_i | e'_i, l) V_n(\mathbf{e}''). \end{aligned}$$

Note that $\bar{V}_{nk}(\mathbf{e}')$ is well-defined if only if firm k is active in state \mathbf{e}' , for if firm k is inactive, then the conditioning event that firm k wins the sale has probability zero. As a convenient normalization, we specify that if firm k is inactive in state \mathbf{e}' , then $\bar{V}_{nk}(\mathbf{e}') = 0$.

The first-order condition that determines an active firm's equilibrium price $p_n(\mathbf{e}')$ can be expressed as:

$$p_n(\mathbf{e}') - c(e'_n) = \frac{\sigma}{1 - D_n(\mathbf{p}(\mathbf{e}'))} - \beta \sum_{k=0}^N \left[\frac{\partial D_k(\mathbf{p}(\mathbf{e}')) / \partial p_n}{-\partial D_n(\mathbf{p}(\mathbf{e}')) / \partial p_n} \right] [\bar{V}_{nn}(\mathbf{e}') - \bar{V}_{nk}(\mathbf{e}')], \quad (\text{B.2})$$

where $\mathbf{p}(\mathbf{e}') = (p_1(\mathbf{e}'), \dots, p_N(\mathbf{e}'))$ is the vector of equilibrium prices in state \mathbf{e}' .

B.3.2 Exit decisions

To develop the Bellman equation determining $V_n(\mathbf{e})$, consider the exit decision $\tau_n(\mathbf{e}, X_n)$ of incumbent firm n who has drawn a random scrap value X_n . An incumbent firm remains in the industry in state \mathbf{e} if its realized scrap value is less than or equal to the expected value of continuing forward to the price-setting stage. Let

$$\hat{X}_n(\mathbf{e}) \equiv E_{\mathbf{e}'}[U_n(\mathbf{e}') | \mathbf{e}, e'_n = e_n, \boldsymbol{\lambda}_{-n}(\mathbf{e})] \quad (\text{B.3})$$

denote the expected value to incumbent firm n of continuing forward to the price-setting stage as an active firm with its current know-how level, i.e., $e'_n = e_n$, taking into account the operating probabilities $\boldsymbol{\lambda}_{-n}(\mathbf{e})$ of all other firms.

$\widehat{X}_n(\mathbf{e})$ is defined as

$$\sum_{e'_1 \in \mathcal{E}'_1} \cdots \sum_{e'_{n-1} \in \mathcal{E}'_{n-1}} \sum_{e'_{n+1} \in \mathcal{E}'_{n+1}} \cdots \sum_{e'_N \in \mathcal{E}'_N} U_n(e'_1, \dots, e'_{n-1}, e_n, e'_{n+1}, \dots, e'_N) \prod_{k \neq n, e'_k \neq 0} \lambda_k(\mathbf{e}) \prod_{k \neq n, e'_k = 0} (1 - \lambda_k(\mathbf{e})),$$

where

$$\mathcal{E}'_n = \begin{cases} \{0, e_n\} & \text{if } e_n \neq 0, \\ \{0, e^0\} & \text{if } e_n = 0. \end{cases}$$

The exit decision of incumbent firm n is characterized by

$$\tau_n(\mathbf{e}, X_n) = \begin{cases} 1 & \text{if } X_n < \widehat{X}_n(\mathbf{e}), \\ [0, 1] & \text{if } X_n = \widehat{X}_n(\mathbf{e}), \\ 0 & \text{if } X_n > \widehat{X}_n(\mathbf{e}). \end{cases}$$

The expected net present value of future cash flows $V_n(\mathbf{e}, X_n)$ to incumbent firm n who has drawn scrap value X_n is given by:

$$V_n(\mathbf{e}, X_n) = \max\{\widehat{X}_n(\mathbf{e}), X_n\}.$$

Integrating over all possible scrap values gives us the value function $V_n(\mathbf{e})$ for incumbent firm n in state \mathbf{e} :

$$\begin{aligned} V_n(\mathbf{e}) &= \int_{\overline{X}-a}^{\overline{X}+a} \max\{\widehat{X}_n(\mathbf{e}), X_n\} dG_X(X_n) \\ &= \begin{cases} \overline{X} & \text{if } \widehat{X}_n(\mathbf{e}) < \overline{X} - a, \\ \frac{1}{4a} [\widehat{X}_n(\mathbf{e})^2 - 2\widehat{X}_n(\mathbf{e})(\overline{X} - a) + (\overline{X} + a)^2] & \text{if } \widehat{X}_n(\mathbf{e}) \in [\overline{X} - a, \overline{X} + a], \\ \widehat{X}_n(\mathbf{e}) & \text{if } \widehat{X}_n(\mathbf{e}) > \overline{X} + a. \end{cases} \end{aligned} \quad (\text{B.4})$$

Since scrap values are private, from the point of view of other firms, the probability $\lambda_n(\mathbf{e})$ that incumbent firm n remains in the industry is:

$$\lambda_n(\mathbf{e}) = G_X(\widehat{X}_n(\mathbf{e})) = \begin{cases} 0 & \text{if } \widehat{X}_n(\mathbf{e}) < \overline{X} - a, \\ \frac{1}{2} + \frac{\widehat{X}_n(\mathbf{e}) - \overline{X}}{2a} & \text{if } \widehat{X}_n(\mathbf{e}) \in [\overline{X} - a, \overline{X} + a], \\ 1 & \text{if } \widehat{X}_n(\mathbf{e}) > \overline{X} + a. \end{cases} \quad (\text{B.5})$$

Because G_X is monotone on the support of X , $\lambda_n(\mathbf{e})$ can be used as a policy choice instead of $\tau_n(\mathbf{e}, X_n)$.

B.3.3 Entry decisions

Consider the entry decision $\tau_n(\mathbf{e}, S_n)$ of potential entrant n who has drawn a random set-up cost S_n . A potential entrant enters the industry in state \mathbf{e} if its realized set-up cost is less than or equal to the expected value of continuing forward to the price-setting stage. Let

$$\widehat{S}_n(\mathbf{e}) \equiv E_{e'}[U_n(\mathbf{e}') | \mathbf{e}, e'_n = e^0, \lambda_{-n}(\mathbf{e})] \quad (\text{B.6})$$

be the expected value of continuing forward to the price-setting stage as an active firm with the initial know-how level, i.e., $e'_n = e^0$, taking into account the operating probabilities $\lambda_{-n}(\mathbf{e})$ of all other firms. $\widehat{S}_n(\mathbf{e})$ is computed analogous to $\widehat{X}_n(\mathbf{e})$, and the entry decision of potential entrant n is characterized by

$$\tau_n(\mathbf{e}, S_n) = \begin{cases} 1 & \text{if } S_n \leq \widehat{S}_n(\mathbf{e}), \\ 0 & \text{if } S_n \geq \widehat{S}_n(\mathbf{e}). \end{cases}$$

The expected net present value of future cash flows $V_n(\mathbf{e}, S_n)$ to potential entrant n who has drawn set-up cost S_n

is given by:

$$V_n(\mathbf{e}, S_n) = \max\{\widehat{S}_n(\mathbf{e}) - S_n, 0\}.$$

Integrating over all possible set-up costs gives us the value function $V_n(\mathbf{e})$ for potential entrant n in state \mathbf{e} :

$$\begin{aligned} V_n(\mathbf{e}) &= \int_{\overline{S}-b}^{\overline{S}+b} \max\{\widehat{S}_n(\mathbf{e}) - S_n, 0\} dG_S(S_n) \\ &= \begin{cases} 0 & \text{if } \widehat{S}_n(\mathbf{e}) < \overline{S}-b, \\ \frac{1}{4b} \left[\widehat{S}_n(\mathbf{e})^2 - 2\widehat{S}_n(\mathbf{e})(\overline{S}-b) + (\overline{S}-b)^2 \right] & \text{if } \widehat{S}_n(\mathbf{e}) \in [\overline{S}-b, \overline{S}+b], \\ \widehat{S}_n(\mathbf{e}) - \overline{S} & \text{if } \widehat{S}_n(\mathbf{e}) > \overline{S}+b. \end{cases} \end{aligned} \quad (\text{B.7})$$

Finally, from the point of view of other firms, the probability $\lambda_n(\mathbf{e})$ that potential entrant n enters the industry is:

$$\lambda_n(\mathbf{e}) = G_S(\widehat{S}_n(\mathbf{e})) = \begin{cases} 0 & \text{if } \widehat{S}_n(\mathbf{e}) < \overline{S}-b, \\ \frac{1}{2} + \frac{\widehat{S}_n - \overline{S}}{2b} & \text{if } \widehat{S}_n(\mathbf{e}) \in [\overline{S}-b, \overline{S}+b], \\ 1 & \text{if } \widehat{S}_n(\mathbf{e}) > \overline{S}+b. \end{cases} \quad (\text{B.8})$$

B.4 Equilibrium

Since our demand specification is symmetric, i.e.,

$$D_n(p_1, \dots, p_{n-1}, p_n, p_{n+1}, \dots, p_N) = D_1(p_n, \dots, p_{n-1}, p_1, p_{n+1}, \dots, p_N)$$

for all n , as well as anonymous, i.e.,

$$D_1(p_1, \dots, p_n, \dots, p_k, \dots, p_N) = D_1(p_1, \dots, p_k, \dots, p_n, \dots, p_N)$$

for all $n \geq 2$ and $k \geq 2$, we focus attention on symmetric and anonymous Markov perfect equilibria (MPE). In such a MPE, it suffices to determine the value and policy functions of a prototypical firm, which we take to be firm 1. We can then obtain the value and policy functions of firm n as follows:

$$\begin{aligned} V_n(\mathbf{e}) &= V_1(\mathbf{e}^{[n]}), \\ p_n(\mathbf{e}) &= p_1(\mathbf{e}^{[n]}), \\ \lambda_n(\mathbf{e}) &= \lambda_1(\mathbf{e}^{[n]}), \end{aligned}$$

where $\mathbf{e}^{[n]} \equiv (e_n, \dots, e_{n-1}, e_1, e_{n+1}, \dots, e_N)$ is the vector formed by switching the first and n th entries in $\mathbf{e} = (e_1, \dots, e_N)$. A symmetric and anonymous MPE is therefore defined by value and policy functions $V_1(\cdot)$, $p_1(\cdot)$, and $\lambda_1(\cdot)$ satisfying equations (B.2), (B.4), (B.5), (B.7), and (B.8). These equations form the basis of the algorithm used to compute the equilibrium.

B.5 Computation

To compute a symmetric and anonymous MPE, we use a variant of the algorithm described in Pakes and McGuire (1994). The algorithm works iteratively. It takes a value function $\widetilde{V}(\mathbf{e})$ and policy functions $\widetilde{p}(\mathbf{e})$ and $\widetilde{\lambda}(\mathbf{e})$ as the starting point for an iteration and generates updated value and policy functions. Each iteration proceeds as follows:

1. Use equation (B.2) to compute an updated pricing policy $p(\mathbf{e})$ for an active firm, taking the pricing policies of all other firms to be given by $\widetilde{p}(\mathbf{e}^{[2]}), \dots, \widetilde{p}(\mathbf{e}^{[N]})$. Use equation (B.1) to compute $U(\mathbf{e})$. In doing so, use $\widetilde{V}(\mathbf{e})$ to compute $\overline{V}_0(\mathbf{e}), \dots, \overline{V}_N(\mathbf{e})$ on the right-hand side of equation (B.2).
2. Use equations (B.5) and (B.8) to compute an updated operating probability $\lambda(\mathbf{e})$ for an incumbent firm and a potential entrant, respectively, using the just-computed $U(\mathbf{e})$ and the operating probabilities $\widetilde{\lambda}(\mathbf{e}^{[2]}), \dots, \widetilde{\lambda}(\mathbf{e}^{[N]})$ for all other firms to calculate the cutoffs $\widehat{X}(\mathbf{e})$ and $\widehat{S}(\mathbf{e})$.

3. Use equations (B.4) and (B.7) to compute the updated value function $V(\mathbf{e})$ for an incumbent firm and a potential entrant, respectively, using the just-computed $U(\mathbf{e})$ and the just-computed operating probability $\lambda(\mathbf{e})$ to calculate the cutoffs $\hat{X}(\mathbf{e})$ and $\hat{S}(\mathbf{e})$.
4. The iteration is completed by assigning $V(\mathbf{e})$ to $\tilde{V}(\mathbf{e})$, $p(\mathbf{e})$ to $\tilde{p}(\mathbf{e})$, and $\lambda(\mathbf{e})$ to $\tilde{\lambda}(\mathbf{e})$.³²

The algorithm terminates once the relative change in the value and the policy functions from one iteration to the next is below a pre-specified tolerance.

B.6 The Cabral-Riordan counterfactual

To implement the Cabral-Riordan “counterfactual” discussed in the text, we begin by noting that in the asymmetric states in the two-firm model, the leader never exited in all of our computed cases in the baseline equilibrium, i.e., $\lambda(e_1, e_2) = 1$ for $e_1 > e_2$. This implies that in identifying possible instances of predatory pricing, we can restrict our attention to cases in which the leader is the low-price firm. This is because if we restrict the follower to act “as if” the leader does not exit in the next period, the follower would behave exactly as it would behave in the baseline model because, indeed, the leader does not exit in the next period.

Given that the leader does not exit in states such that $e_1 > e_2$, the leader’s value in these states is given by

$$V(e_1, e_2) = \lambda(e_2, e_1)U(e_1, e_2) + (1 - \lambda(e_2, e_1))U(e_1, 0),$$

where $U(e_1, 0)$ is the value that the leader gets if the follower exits and $1 - \lambda(e_2, e_1)$ is the probability that the follower exits. In the counterfactual analysis, we imagine that in states $e_1 > e_2$ an industry leader acts “as if” $\lambda(e_2, e_1) = 1$, and so it “believes” $V(e_1, e_2) = U(e_1, e_2)$. This implies that when we compute the MPE in the counterfactual, the expected value functions $\bar{V}_k(e_1, e_2)$ in the asymmetric states ($e_1 > e_2$) will be given by

$$\bar{V}_k(e_1, e_2) = E_{e'_1, e'_2} \{U(e_1, e_2) | (e_1, e_2), \text{ firm } k \text{ wins}\},$$

rather than by

$$\bar{V}_k(e_1, e_2) = E_{e'_1, e'_2} \{V(e_1, e_2) | (e_1, e_2), \text{ firm } k \text{ wins}\}.$$

In all other respects, the computation of the MPE proceeds as described above (specialized to $N = 2$). Note that the counterfactual analysis does not prevent the follower from exiting. Rather, the leader acts “as if” the follower cannot exit in the next period and thus in the leader’s pricing decision, the leader does not take into account the “prize” that arises from the fact that the follower might exit in the next period. Due to this, the counterfactual model loses the rational expectations property: the follower’s actions do not necessarily match the leader’s assumptions on them. In the counterfactual exercise, we take it that the leader cannot update these assumptions by observing the actions of the follower.

References

- ALCHIAN, A. (1963): “Reliability of progress curves in airframe production,” *Econometrica*, 31(4), 679–693.
- ARGOTE, L., AND D. EPPLE (1990): “Learning curves in manufacturing,” *Science*, 247(4945), 920–924.
- BAILEY, CHARLES, D. (1989): “Forgetting and the learning curve,” *Management Science*, 35(3), 340–352.
- BAIN, J. S. (1956): *Barriers to new competition*. Harvard University Press, Cambridge.
- BENKARD, L. (2000): “Learning and forgetting: The dynamics of aircraft production,” *American Economic Review*, 90(4), 1034–1054.
- (2003): “A dynamic analysis of the market for wide-bodied commercial aircraft,” *Review of Economic Studies*, forthcoming.

³²More precisely, we assign a weighted average of $V(\mathbf{e})$ and $\tilde{V}(\mathbf{e})$ to $\tilde{V}(\mathbf{e})$ to help the algorithm converge.

- CABRAL, L., AND M. RIORDAN (1994): “The learning curve, market dominance, and predatory pricing,” *Econometrica*, 62(2), 1115–1140.
- (1997): “The learning curve, predation, antitrust, and welfare,” *Journal of Industrial Economics*, (2), 155–169.
- DARR, E., L. ARGOTE, AND D. EPPLE (1995): “The Acquisition, Transfer, and Depreciation of Knowledge in Service Organizations: Productivity in Franchises,” *Management Science*, 41(11), 1750–1762.
- DORASZELSKI, U., AND M. SATTERTHWAITE (2004): “Foundations of Markov-perfect industry dynamics: Existence, purification, and multiplicity,” Working paper, Hoover Institution, Stanford.
- DUTTON, J., AND A. THOMAS (1984): “Treating progress functions as a managerial opportunity,” 9(2), 235–247.
- ERICSON, R., AND A. PAKES (1995): “Markov-perfect industry dynamics: A framework for empirical work,” *Review of Economic Studies*, 62, 53–82.
- FUDENBERG, D., AND J. TIROLE (1983): “Learning-by-doing and market performance,” *Bell Journal of Economics*, 14(2), 522–530.
- GHEMAWAT, P., AND M. SPENCE (1985): “Learning Curve Spillovers and Market Performance,” *Quarterly Journal of Economics*, 100(Supplement), 839–852.
- HERINGS, J., AND R. PEETERS (2004): “Stationary equilibria in stochastic games: Structure, selection, and computation,” *Journal of Economic Theory*, forthcoming.
- IRWIN, D., AND P. KLENOW (1994): “Learning-by-doing spillovers in the semiconductor industry,” *Journal of Political Economy*, 102(6), 1200–1227.
- LEWIS, T., AND H. YILDIRIM (2002): “Managing Dynamic Competition,” *American Economic Review*, 92(4), 779–797.
- LIEBERMAN, M. (1984): “The Learning Curve and Pricing in the Chemical Processing Industries,” *Rand Journal of Economics*, 15, 213–228.
- ORDOVER, J. A., AND R. D. WILLG (1981): “An Economic Definition of Predation: Pricing and Product Innovation,” *Yale Law Journal*, 91, 8–53.
- PAKES, A., AND P. MCGUIRE (1994): “Computing Markov-perfect Nash equilibria: Numerical implications of a dynamic differentiated product model,” *Rand Journal of Economics*, 25(4), 555–589.
- ROSS, D. (1986): “Learning to Dominate,” *Journal of Industrial Economics*, 34, 337–353.
- SPENCE, M. (1981): “The Learning Curve and Competition,” *Bell Journal of Economics*, 12, 49–70.
- THOMPSON, P. (2001): “How much did the Liberty shipbuilders learn? New evidence for an old case study,” *Journal of Political Economy*, 109, 103–137.
- (2003): “How much did the Liberty shipbuilders forget?,” Working paper, Florida International University, Florida International University.
- THORNTON, R., AND P. THOMPSON (2001): “Learning from Experience and Learning from Others: An Exploration of Learning and Spillovers in Wartime Shipbuilding,” *American Economic Review*, 91, 1350–1368.
- VIVES, X. (1999): *Oligopoly pricing: old ideas and new tools*. MIT Press, Cambridge.
- WIXTED, M., AND E. EBBESEN (1991): “On the form of forgetting,” *Psychological Science*, 2, 409–415.
- ZIMMERMAN, M. (1982): “Learning Effects and the Commercialization of New Energy Technologies: The Case of Nuclear Power,” *Bell Journal of Economics*, 13, 297–310.