

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

David Besanko[†] Ulrich Doraszelski[‡] Yaroslav Kryukov[§]
Mark Satterthwaite[¶]

December 21, 2008

Abstract

Learning-by-doing and organizational forgetting have been shown to be important in a variety of industrial settings. This paper provides a general model of dynamic competition that accounts for these economic fundamentals and shows how they shape industry structure and dynamics. Previously obtained results regarding the dominance properties of firms' pricing behavior no longer hold in this more general setting. We show that forgetting does not simply negate learning. Rather, learning and forgetting are distinct economic forces. In particular, a model with learning and forgetting can give rise to aggressive pricing behavior, market dominance, and multiple equilibria, whereas a model with learning alone cannot.

*We have greatly benefitted from the comments and suggestions of the Editor, Steve Berry, and two anonymous referees. We are also indebted to Lanier Benkard, Luis Cabral, Jiawei Chen, Stefano Demichelis, Michaela Draganska, Ken Judd, Pedro Marin, Ariel Pakes, Michael Ryall, Karl Schmedders, Chris Shannon, Kenneth Simons, Scott Stern, Mike Whinston, and Huseyin Yildirim as well as the participants of various conferences. Guy Arie and Paul Grieco provided excellent research assistance. Besanko and Doraszelski gratefully acknowledge financial support from the National Science Foundation under Grant No. 0615615. Doraszelski further benefitted from the hospitality of the Hoover Institution during the academic year 2006/07. Kryukov thanks the General Motors Center for Strategy in Management at Northwestern's Kellogg School of Management for support during this project. Satterthwaite acknowledges gratefully that this material is based upon work supported by the National Science Foundation under Grant No. 0121541.

[†]Kellogg School of Management, Northwestern University, Evanston, IL 60208, d-besanko@kellogg.northwestern.edu.

[‡]Department of Economics, Harvard University, Cambridge, MA 02138, doraszelski@harvard.edu.

[§]Department of Economics, Northwestern University, Evanston, IL 60208, kryukov@northwestern.edu.

[¶]Kellogg School of Management, Northwestern University, Evanston, IL 60208, m-satterthwaite@kellogg.northwestern.edu.

1 Introduction

Empirical studies provide ample evidence that the marginal cost of production decreases with cumulative experience in a variety of industrial settings. This fall in marginal cost is known as learning-by-doing. More recent empirical studies also suggest that organizations can forget the know-how gained through learning-by-doing due to labor turnover, periods of inactivity, and failure to institutionalize tacit knowledge.¹ Organizational forgetting has been largely ignored by the theoretical literature. This is problematic because Benkard (2004) shows that organizational forgetting is essential to explain the dynamics in the market for wide-bodied airframes in the 1970s and 1980s.

In this paper we build on the computational Markov-perfect equilibrium framework of Ericson & Pakes (1995) to analyze how the economic fundamentals of learning-by-doing and organizational forgetting interact to determine industry structure and dynamics.² We account for organizational forgetting in Cabral & Riordan's (1994) (C-R) seminal model of learning-by-doing.³ This seemingly small change has surprisingly large effects. Dynamic competition with learning and forgetting is akin to racing down an upward moving escalator. As long as a firm makes sales sufficiently frequently so that the gain in know-how from learning outstrips the loss in know-how from forgetting, it moves down its learning curve and its marginal cost decreases. However, if sales slow down or come to a halt, perhaps because of its competitors' aggressive pricing, then the firm slides back up its learning curve and its marginal cost increases. This cannot happen in the C-R model. Due to this qualitative difference, organizational forgetting leads to a rich array of pricing behaviors and industry dynamics that the existing literature neither imagined nor explained.

It is often said that learning-by-doing promotes market dominance because it gives a more experienced firm the ability to profitably underprice its less experienced rival and therefore shut out the competition in the long run. As Dasgupta & Stiglitz (1988) explain:

... firm-specific learning encourages the growth of industrial concentration.

To be specific, one expects that strong learning possibilities, coupled with vigorous competition among rivals, ensures that history matters ... in the sense that

¹See Wright (1936)HIRS:52, DeJong (1957), Alchian (1963), Levy (1965), Kilbridge (1962), Hirschmann (1964), Preston & Keachie (1964), Baloff (1971), Dudley (1972), Zimmerman (1982), Lieberman (1984), Gruber (1992), Irwin & Klenow (1994), Jarmin (1994), Pisano (1994), Bohn (1995), Hatch & Mowery (1998), Thompson (2001), and Thornton & Thompson (2001) for empirical studies of learning-by-doing and Argote, Beckman & Epple (1990), Darr, Argote & Epple (1995), Benkard (2000), Shafer, Nembhard & Uzumeri (2001), and Thompson (2003) for organizational forgetting.

²Dynamic stochastic games and feedback strategies that map states into actions date back at least to Shapley (1953). Maskin & Tirole (2001) provide the fundamental theory showing how many subgame perfect equilibria of these games can be represented consistently and robustly as Markov perfect equilibria.

³Prior to the infinite-horizon price-setting model of C-R, the literature had studied learning-by-doing using finite-horizon quantity-setting models (Spence 1981, Fudenberg & Tirole 1983, Ghemawat & Spence 1985, Ross 1986, Dasgupta & Stiglitz 1988, Cabral & Riordan 1997).

if a given firm enjoys some initial advantages over its rivals it can, by undercutting them, capitalize on these advantages in such a way that the advantages accumulate over time, rendering rivals incapable of offering effective competition in the long run ... (p. 247)

However, if organizational forgetting “undoes” learning-by-doing, then forgetting may be a procompetitive force and an antidote to market dominance through learning. Two reasons for suspecting this come to mind. First, to the extent that the leader has more to forget than the follower, forgetting should work to equalize differences between firms. Second, because forgetting makes improvements in competitive position from learning transitory, it should make firms reluctant to invest in the acquisition of know-how through price cuts. We reach the opposite conclusion: Organizational forgetting can make firms more aggressive rather than less aggressive. This aggressive pricing behavior, in turn, puts the industry on a path towards market dominance.

In the absence of organizational forgetting, the price that a firm sets reflects two goals. First, by winning a sale, the firm moves down its learning curve. This is the advantage-building motive. Second, the firm prevents its rival from moving down its learning curve. This is the advantage-defending motive. But in the presence of organizational forgetting bidirectional movements through the state space are possible, and this opens up new strategic possibilities for building and defending advantage. By winning a sale, a firm makes itself less vulnerable to forgetting by creating a “buffer stock” of know-how. By winning and denying its rival a buffer stock, the firm also makes its rival more vulnerable to forgetting. Because organizational forgetting reinforces the advantage-building and advantage-defending motives in this way, it can create strong incentives to cut prices in order to win the sale. Organizational forgetting is thus a source of aggressive pricing behavior.

While the existing literature has mainly focused on the dominance properties of firms’ pricing behaviors, we find that these properties are neither necessary nor sufficient for market dominance in our more general setting. We therefore go beyond the existing literature and directly examine the industry dynamics that firms’ pricing behaviors imply. We find that organizational forgetting is a source of—not an antidote to—market dominance. If forgetting is sufficiently weak, then asymmetries may arise but cannot persist. As in the C-R model, learning-by-doing operates as a ratchet: firms inexorably—if at different rates—move towards the bottom of their learning curves where cost parity obtains. If forgetting is sufficiently strong, then asymmetries cannot arise in the first place because forgetting stifles investment in learning altogether. But, for intermediate degrees of forgetting, asymmetries arise and persist. Even extreme asymmetries akin to near-monopoly are possible. This is because, in the presence of organizational forgetting, the leader can use price cuts to delay or even stall the follower in moving down its learning curve.

Organizational forgetting is also a source of multiple equilibria. If the inflow of know-how into the industry due to learning is substantially smaller than the outflow of know-how due to forgetting, then it is virtually impossible that both firms reach the bottom of their learning curves. Conversely, if the inflow is substantially greater than the outflow, then it is virtually inevitable that they do reach the bottom. In both cases, the primitives of the model tie down the equilibrium. This is no longer the case if the inflow roughly balances the outflow. If both firms believe that they cannot profitably coexist at the bottom of their learning curves, then both cut their prices in the hope of acquiring a competitive advantage early on and maintaining it throughout. However, if both firms believe that they can profitably coexist, then neither cuts its price, thereby ensuring that the anticipated symmetric industry structure emerges. Consequently, in addition to the degree of forgetting, the equilibrium by itself is an important determinant of pricing behavior and industry dynamics.

Our finding of multiplicity is important for two reasons. First, to our knowledge, all applications of Ericson & Pakes's (1995) framework have found a single equilibrium. Pakes & McGuire (1994) (P-M) indeed held that "nonuniqueness does not seem to be a problem" (p. 570). It is therefore striking that we obtain up to nine equilibria for some parameterizations. Second, we show that multiple equilibria in our model arise from firms' expectations regarding the value of continued play. Being able to pinpoint the driving force behind multiple equilibria is a first step towards tackling the multiplicity problem that plagues the estimation of dynamic stochastic games and inhibits the use of counterfactuals in policy analysis.⁴

In sum, learning-by-doing and organizational forgetting are distinct economic forces. Forgetting, in particular, does not simply negate learning. The unique role forgetting plays comes about because it makes bidirectional movements through the state space possible. Thus the interaction of learning and forgetting can give rise to aggressive pricing behavior, long-run market dominance of varying degrees, and multiple equilibria.

We emphasize that organizational forgetting is not unique in leading to long-run market dominance. As we show in Section 8, a model of learning-by-doing that incorporates shut-out elements such as a choke price or entry and exit may also do so, much in the way Dasgupta & Stiglitz (1988) describe. Following C-R we exclude shut-out elements from our basic model for two reasons. First, the interaction between learning and forgetting is subtle and generates an enormous variety of interesting, even surprising, equilibria. Isolating it is therefore useful theoretically. Second, from an empirical viewpoint, as with Intel and AMD we may see an apparently stable hierarchy of firms with differing market shares, costs, and profits. Our model can generate such an outcome, while shut-out model elements favor more extreme outcomes where either one firm dominates or all firms compete on equal footing.

⁴See Akerberg, Benkard, Berry & Pakes (2007) and Pakes (2008) for a discussion of the issue.

We also make two methodological contributions. First, we point out a weakness of the P-M algorithm, the major tool for computing equilibria in the literature following Ericson & Pakes (1995). Specifically, we prove that our dynamic stochastic game has equilibria that the P-M algorithm cannot compute. Roughly speaking, in the presence of multiple equilibria, “in between” two equilibria that it can compute there is one equilibrium it cannot. This severely limits its ability to provide a complete picture of the set of solutions to the model.

Second, we propose a homotopy or path-following algorithm. The algorithm traces out the equilibrium correspondence by varying the degree of forgetting. This allows us to compute equilibria that the P-M algorithm cannot compute. We find that the equilibrium correspondence contains a unique path that starts at the equilibrium of the C-R model. Whenever this path bends back on itself and then forward again, there are multiple equilibria. In addition, the equilibrium correspondence may contain one or more loops that cause additional multiplicity. To our knowledge, our paper is the first to describe in detail the structure of the set of equilibria of a dynamic stochastic game in the tradition of Ericson & Pakes (1995).

The organization of the remainder of the paper is as follows. Sections 2 and 3 describe the model specification and our computational strategy. Section 4 provides an overview of the equilibrium correspondence. Section 5 analyzes industry dynamics and Section 6 characterizes the pricing behavior that drives it. Section 7 describes how organizational forgetting can lead to multiple equilibria. Section 8 undertakes a number of robustness checks. Section 9 summarizes and concludes.

Throughout the paper in presenting our findings we distinguish between *results*, which are established numerically through a systematic exploration of a subset of the parameter space, and *propositions*, which hold true for the entire parameter space. If a proposition establishes a possibility through an example, then the example is presented adjacent to the proposition. If the proof of a proposition is deductive, then it is contained in the Appendix.

2 Model

For expositional clarity we focus on the basic model of an industry with two firms and neither entry nor exit; the Online Appendix outlines the general model. Our basic model is the C-R model with organizational forgetting added and, to allow for our computational approach, specific functional forms for demand and cost.

Firms and states. We consider a discrete-time, infinite-horizon dynamic stochastic game of complete information played by two firms. Firm $n \in \{1, 2\}$ is described by its state

$e_n \in \{1, \dots, M\}$. A firm's state indicates its cumulative experience or stock of know-how. By making a sale, a firm can add to its stock of know-how. Following C-R, we use a period just long enough for a firm to make a sale.⁵ As suggested by the empirical studies of Argote et al. (1990), Darr et al. (1995), Benkard (2000), Shafer et al. (2001), and Thompson (2003) we account for organizational forgetting. Accordingly, the evolution of firm n 's stock of know-how is governed by the law of motion

$$e'_n = e_n + q_n - f_n,$$

where e'_n and e_n is firm n 's stock of know-how in the subsequent and current period, respectively, the random variable $q_n \in \{0, 1\}$ indicates whether firm n makes a sale and gains a unit of know-how through learning-by-doing, and the random variable $f_n \in \{0, 1\}$ indicates whether firm n loses a unit of know-how through organizational forgetting.

At any point in time, the industry is characterized by a vector of firms' states $\mathbf{e} = (e_1, e_2) \in \{1, \dots, M\}^2$. We refer to \mathbf{e} as the state of the industry. We use $\mathbf{e}^{[2]}$ to denote the vector (e_2, e_1) constructed by interchanging the stocks of know-how of firms 1 and 2.

Learning-by-doing. Firm n 's marginal cost of production $c(e_n)$ depends on its stock of know-how e_n through a learning curve

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

where $\eta = \log_2 \rho$ for a progress ratio of $\rho \in (0, 1]$. Marginal cost decreases by $100(1 - \rho)$ percent as the stock of know-how doubles, so that a lower progress ratio implies a steeper learning curve. The marginal cost of production at the top of the learning curve, $c(1)$, is $\kappa > 0$ and, as in C-R, m represents the stock of know-how at which a firm reaches the bottom of its learning curve.⁶

Organizational forgetting. We let $\Delta(e_n) = \Pr(f_n = 1)$ denote the probability that firm n loses a unit of know-how through organizational forgetting. We assume that this probability is nondecreasing in the firm's experience level. This has several advantages. First, experimental evidence in the management literature suggests that forgetting by individuals is an increasing function of the current stock of learned knowledge (Bailey 1989).

⁵A sale may involve a single unit or a batch of units (e.g., 100 aircraft or 10,000 memory chips) that are sold to a single buyer.

⁶While C-R take the state space to be infinite, i.e., $M = \infty$ in our notation, they make the additional assumption that the price that a firm charges does not depend on how far it is beyond the bottom of its learning curve (p. 1119). This is tantamount to assuming, as we do, that the state space is finite.

Second, a direct implication of $\Delta(\cdot)$ increasing is that the expected stock of know-how in the absence of further learning is a decreasing convex function of time.⁷ This phenomenon, known in the psychology literature as Jost’s second law, is consistent with experimental evidence on forgetting by individuals (Wixted & Ebbesen 1991). Third, in the capital-stock model employed in empirical work on organizational forgetting, the amount of depreciation is assumed to be proportional to the stock of know-how. Hence, the additional know-how needed to counteract depreciation must increase with the stock of know-how. Our specification has this feature. However, unlike the capital-stock model, it is consistent with a discrete state space.⁸

The specific functional form we employ is

$$\Delta(e_n) = 1 - (1 - \delta)^{e_n},$$

where $\delta \in [0, 1]$ is the forgetting rate.⁹ If $\delta > 0$, then $\Delta(e_n)$ is increasing and concave in e_n ; $\delta = 0$ corresponds to the absence of organizational forgetting, the special case C-R analyzed. Other functional forms are plausible, and we explore one of them in the Online Appendix.

Demand. The industry draws its customers from a large pool of potential buyers. In each period, one buyer enters the market and purchases the good from one of the two firms. The utility that the buyer obtains by purchasing good n is $v - p_n + \varepsilon_n$ where p_n is the price of good n , v is a deterministic component of utility, and ε_n is a stochastic component that captures the idiosyncratic preference for good n of this period’s buyer. Both ε_1 and ε_2 are unobservable to firms and are independently and identically type 1 extreme value distributed with location parameter 0 and scale parameter $\sigma > 0$. The scale parameter governs the degree of horizontal product differentiation. As $\sigma \rightarrow 0$, goods become homogeneous.

The buyer purchases the good that gives it the highest utility. Given our distributional assumptions the probability that firm n makes the sale is given by the logit specification

$$D_n(\mathbf{p}) = \Pr(q_n = 1) = \frac{\exp(\frac{v-p_n}{\sigma})}{\sum_{k=1}^2 \exp(\frac{v-p_k}{\sigma})} = \frac{1}{1 + \exp(\frac{p_n-p_{-n}}{\sigma})},$$

where $\mathbf{p} = (p_1, p_2)$ is the vector of prices and p_{-n} denotes the price the other firm charges. Demand effectively depends on differences in prices because we assume, as do C-R, that the

⁷See the Online Appendix for a proof.

⁸See Benkard (2004) for an alternative approximation to the capital-stock model.

⁹One way to motivate this functional form is to imagine that the stock of know-how is dispersed among a firm’s workforce. In particular, assume that e_n is the number of skilled workers and that organizational forgetting is the result of labor turnover. Then, given a turnover rate of δ , $\Delta(e_n)$ is the probability that at least one of the e_n skilled workers leaves the firm.

buyer always purchases from one of the two firms in the industry. In Section 8 we discuss the effects of including an outside good in the specification.

State-to-state transitions. From one period to the next, a firm's stock of know-how moves up or down or remains constant depending on realized demand $q_n \in \{0, 1\}$ and organizational forgetting $f_n \in \{0, 1\}$. The transition probabilities are

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

where, at the upper and lower boundaries of the state space, we modify the transition probabilities to be $\Pr(M|M, 1) = 1$ and $\Pr(1|1, 0) = 1$, respectively. Note that a firm can increase its stock of know-how only if it makes a sale in the current period, an event that has probability $D_n(\mathbf{e})$; otherwise it runs the risk that its stock of know-how decreases.

Bellman equation. We define $V_n(\mathbf{e})$ to be the expected net present value of firm n 's cash flows if the industry is currently in state \mathbf{e} . The value function $\mathbf{V}_n : \{1, \dots, M\}^2 \rightarrow [-\hat{V}, \hat{V}]$, where \hat{V} is a sufficiently large constant, is implicitly defined by the Bellman equation

$$V_n(\mathbf{e}) = \max_{p_n} D_n(p_n, p_{-n}(\mathbf{e})) (p_n - c(e_n)) + \beta \sum_{k=1}^2 D_k(p_n, p_{-n}(\mathbf{e})) \bar{V}_{nk}(\mathbf{e}), \quad (1)$$

where $p_{-n}(\mathbf{e})$ is the price charged by the other firm in state \mathbf{e} , $\beta \in (0, 1)$ is the discount factor, and $\bar{V}_{nk}(\mathbf{e})$ is the expectation of firm n 's value function conditional on the buyer purchasing the good from firm $k \in \{1, 2\}$ in state \mathbf{e} as given by

$$\bar{V}_{n1}(\mathbf{e}) = \sum_{e'_1=e_1}^{e_1+1} \sum_{e'_2=e_2-1}^{e_2} V_n(\mathbf{e}') \Pr(e'_1 | e_1, 1) \Pr(e'_2 | e_2, 0), \quad (2)$$

$$\bar{V}_{n2}(\mathbf{e}) = \sum_{e'_1=e_1-1}^{e_1} \sum_{e'_2=e_2}^{e_2+1} V_n(\mathbf{e}') \Pr(e'_1 | e_1, 0) \Pr(e'_2 | e_2, 1). \quad (3)$$

The policy function $\mathbf{p}_n : \{1, \dots, M\}^2 \rightarrow [-\hat{p}, \hat{p}]$, where \hat{p} is a sufficiently large constant, specifies the price $p_n(\mathbf{e})$ that firm n sets in state \mathbf{e} .¹⁰ Let $h_n(\mathbf{e}, p_n, p_{-n}(\mathbf{e}), \mathbf{V}_n)$ denote the maximand in the Bellman equation (1). Differentiating it with respect to p_n and using the properties of logit demand we obtain the first-order condition (FOC):

$$0 = \frac{\partial h_n(\cdot)}{\partial p_n} = \frac{1}{\sigma} D_n(p_n, p_{-n}(\mathbf{e})) \left(\sigma - (p_n - c(e_n)) - \beta \bar{V}_{nn}(\mathbf{e}) + h_n(\cdot) \right).$$

¹⁰In what follows we assume that \hat{p} is chosen large enough to not constrain pricing behavior.

Differentiating $h_n(\cdot)$ a second time yields

$$\frac{\partial^2 h_n(\cdot)}{\partial p_n^2} = \frac{1}{\sigma} \frac{\partial h_n(\cdot)}{\partial p_n} \left(2D_n(p_n, p_{-n}(\mathbf{e})) - 1 \right) - \frac{1}{\sigma} D_n(p_n, p_{-n}(\mathbf{e})).$$

If the FOC is satisfied, then $\frac{\partial^2 h_n(\cdot)}{\partial p_n^2} = -\frac{1}{\sigma} D_n(p_n, p_{-n}(\mathbf{e})) < 0$. $h_n(\cdot)$ is therefore strictly quasi-concave in p_n , so that the pricing decision $p_n(\mathbf{e})$ is uniquely determined by the solution to the FOC (given $p_{-n}(\mathbf{e})$).

Equilibrium. In our model firms face identical demand and cost primitives. Asymmetries between firms arise endogenously from the effects of their pricing decisions on realized demand and organizational forgetting. Hence, we focus attention on symmetric Markov perfect equilibria (MPE). In a symmetric equilibrium the pricing decision taken by firm 2 in state \mathbf{e} is identical to the pricing decision taken by firm 1 in state $\mathbf{e}^{[2]}$, i.e., $p_2(\mathbf{e}) = p_1(\mathbf{e}^{[2]})$, and similarly for the value function. It therefore suffices to determine the value and policy functions of firm 1. We define $V(\mathbf{e}) = V_1(\mathbf{e})$ and $p(\mathbf{e}) = p_1(\mathbf{e})$ for each state \mathbf{e} . Further, we let $\bar{V}_k(\mathbf{e}) = \bar{V}_{1k}(\mathbf{e})$ denote the conditional expectation of firm 1's value function and $D_k(\mathbf{e}) = D_k(p(\mathbf{e}), p(\mathbf{e}^{[2]}))$ denote the probability that the buyer purchases from firm $k \in \{1, 2\}$ in state \mathbf{e} .

Given this notation, the Bellman equation and FOC can be expressed as

$$F_{\mathbf{e}}^1(\mathbf{V}^*, \mathbf{p}^*) = -V^*(\mathbf{e}) + D_1^*(\mathbf{e})(p^*(\mathbf{e}) - c(e_1)) + \beta \sum_{k=1}^2 D_k^*(\mathbf{e}) \bar{V}_k^*(\mathbf{e}) = 0, \quad (4)$$

$$F_{\mathbf{e}}^2(\mathbf{V}^*, \mathbf{p}^*) = \sigma - (1 - D_1^*(\mathbf{e}))(p^*(\mathbf{e}) - c(e_1)) - \beta \bar{V}_1^*(\mathbf{e}) + \beta \sum_{k=1}^2 D_k^*(\mathbf{e}) \bar{V}_k^*(\mathbf{e}) = 0, \quad (5)$$

where we use asterisks to denote an equilibrium. The collection of equations (4) and (5) for all states $\mathbf{e} \in \{1, \dots, M\}^2$ can be written more compactly as

$$\mathbf{F}(\mathbf{V}^*, \mathbf{p}^*) = \begin{bmatrix} F_{(1,1)}^1(\mathbf{V}^*, \mathbf{p}^*) \\ F_{(2,1)}^1(\mathbf{V}^*, \mathbf{p}^*) \\ \vdots \\ F_{(M,M)}^2(\mathbf{V}^*, \mathbf{p}^*) \end{bmatrix} = \mathbf{0}, \quad (6)$$

where $\mathbf{0}$ is a $(2M^2 \times 1)$ vector of zeros. Any solution to this system of $2M^2$ equations in $2M^2$ unknowns $\mathbf{V}^* = (V^*(1, 1), V^*(2, 1), \dots, V^*(M, M))$ and $\mathbf{p}^* = (p^*(1, 1), p^*(2, 1), \dots, p^*(M, M))$ is a symmetric equilibrium in pure strategies. A slightly modified version of Proposition 2 in Doraszelski & Satterthwaite (2008) establishes that such an equilibrium always exists.

Baseline parameterization. Since our focus is on how learning-by-doing and organizational forgetting affect pricing behavior, and the industry dynamics this behavior implies, we explore the full range of values for the progress ratio ρ and the forgetting rate δ . To do so, we fix the remaining parameters to their baseline values given below. We specify a grid of 100 equidistant values of $\rho \in (0, 1]$. For each of them, we use the homotopy algorithm described in Section 3 to trace the equilibrium as δ ranges from 0 to 1. Typically this entails solving the model for a few thousand intermediate values of δ . If an important or interesting property is true for each of these systematically computed equilibria, then we report it as a result. In Section 8 we then vary the values of the parameters other than ρ and δ in order to discuss their influence on the equilibrium and demonstrate the robustness of our conclusions.

While we explore the full range of values for ρ and δ , we note that most empirical estimates of progress ratios are in the range of 0.7 to 0.95 (Dutton & Thomas 1984). However, a very steep learning curve, with ρ much less than 0.7, may also capture a practically relevant situation. Suppose the first unit of a product is a hand-built prototype and the second unit is a guinea pig for organizing the production line. After this point the gains from learning-by-doing are more or less exhausted and the marginal cost of production is close to zero.¹¹ Benkard (2000) and Argote et al. (1990) have found monthly rates of depreciation ranging from 4 to 25 percent of the stock of know-how. In the Online Appendix we show how to map these estimates that are based on a capital-stock model of organizational forgetting into in our specification. The implied values of the forgetting rate δ fall below 0.1.

In our baseline parameterization we set $M = 30$ and $m = 15$. The marginal cost at the top of the learning curve κ is equal to 10. For a progress ratio of $\rho = 0.85$, this implies that the marginal cost of production declines from a maximum value of $c(1) = 10$ to a minimum value of $c(15) = \dots = c(30) = 5.30$. For $\rho = 0.15$, we have the case of a hand-built prototype where the marginal cost of production declines very quickly from $c(1) = 10$ over $c(2) = 1.50$ and $c(3) = 0.49$ to $c(15) = \dots = c(30) = 0.01$.

Turning to demand, we set $\sigma = 1$ in our baseline parameterization. To illustrate, in the Nash equilibrium of a static price-setting game (obtained by setting $\beta = 0$ in our model), the own-price elasticity of demand ranges between -8.86 in state $(1, 15)$ and -2.13 in state $(15, 1)$ for a progress ratio of $\rho = 0.85$. The cross-price elasticity of firm 1's demand with respect to firm 2's price is 2.41 in state $(15, 1)$ and 7.84 in state $(1, 15)$. For $\rho = 0.15$ the own-price elasticity ranges between -9.89 and -1.00 and the cross-price elasticity between 1.00 and 8.05. These reasonable elasticities suggest that the results reported below are not artifacts of extreme parameterizations.

¹¹To avoid a marginal cost of close to zero, shift the cost function $c(e_n)$ by $\tau > 0$. While introducing a component of marginal cost that is unresponsive to learning-by-doing shifts the policy function by τ , the value function and the industry dynamics remain unchanged.

We finally set the discount factor to $\beta = \frac{1}{1.05}$. It may be thought of as $\beta = \frac{\zeta}{1+r}$, where $r > 0$ is the per-period discount rate and $\zeta \in (0, 1]$ is the exogenous probability that the industry survives from one period to the next. Consequently, our baseline parameterization corresponds to a variety of scenarios that differ in the length of a period. For example, it corresponds to a period length of one year, a yearly discount rate of 5 percent, and certain survival. Perhaps more interestingly, it also corresponds to a period length of one month, a monthly discount rate of 1 percent (which translates into a 12.68 percent annual discount rate), and a monthly survival probability of 0.96. To put this—our focal scenario—in perspective, technology companies such as IBM and Microsoft had costs of capital in the range of 11 to 15 percent per annum in the late 1990s. Further, an industry with a monthly survival probability of 0.96 has an expected lifetime of 26.25 months. Thus this scenario is consistent with a pace of innovative activity that is expected to make the current generation of products obsolete within two to three years.

3 Computation

In this section we first describe a novel algorithm for computing equilibria of dynamic stochastic games that is based on the homotopy method.¹² Then, we turn to the P-M algorithm that is the standard means for computing equilibria in the literature following Ericson & Pakes (1995). We show that it is inadequate for characterizing the set of solutions to our model although it remains useful for obtaining a starting point for the homotopy algorithm. A reader who is more interested in the economic implications of learning and forgetting may skip ahead to Section 4 after reading the first part of this section that introduces the homotopy algorithm by way of an example.

3.1 Homotopy algorithm

Our goal is to explore the graph of the equilibrium correspondence as the forgetting rate δ and the progress ratio ρ vary:

$$\mathbf{F}^{-1} = \{(\mathbf{V}^*, \mathbf{p}^*, \delta, \rho) | \mathbf{F}(\mathbf{V}^*, \mathbf{p}^*; \delta, \rho) = \mathbf{0}, \delta \in [0, 1], \rho \in (0, 1]\}, \quad (7)$$

where $\mathbf{F}(\cdot)$ is the system of equations (6) that defines an equilibrium and we make explicit that it depends on δ and ρ (recall that we hold fixed the remaining parameters). The graph \mathbf{F}^{-1} is a surface, or set of surfaces, that may have folds. Our homotopy algorithm explores

¹²See Schmedders (1998, 1999) for an application of the homotopy method to general equilibrium models with incomplete asset markets and Berry & Pakes (2007) for an application to estimating demand systems.

this graph by taking slices of it for given values of ρ :

$$\mathbf{F}^{-1}(\rho) = \{(\mathbf{V}^*, \mathbf{p}^*, \delta) | \mathbf{F}(\mathbf{V}^*, \mathbf{p}^*; \delta, \rho) = \mathbf{0}, \delta \in [0, 1]\}. \quad (8)$$

The homotopy algorithm starts from a single equilibrium that has already been computed and traces out an entire path of equilibria in $\mathbf{F}^{-1}(\rho)$ by varying δ . The homotopy algorithm is therefore also called a path-following algorithm and δ the homotopy parameter.

Example. An example helps explain how the homotopy algorithm works. Consider the equation $F(x; \delta) = 0$, where

$$F(x; \delta) = -15.289 - \frac{\delta}{1 + \delta^4} + 67.500x - 96.923x^2 + 46.154x^3. \quad (9)$$

Equation (9) implicitly relates an endogenous variable x to an exogenous parameter δ . Figure 1 graphs the set of solutions $F^{-1} = \{(x, \delta) | F(x; \delta) = 0, \delta \in [0, 1]\}$. There are multiple solutions at $\delta = 0.3$, namely $x = 0.610$, $x = 0.707$, and $x = 0.783$. Finding them is trivial with the graph in hand, but even for this simple case the graph is less than straightforward to draw. Whether one solves $F(x; \delta) = 0$ for x taking δ as given or for δ taking x as given, the result is a correspondence, not a function.

The homotopy method introduces an auxiliary variable s that indexes each point on the graph, starting at point A for $s = 0$ and ending at point D for $s = \bar{s}$. The graph is then just the parametric path given by a pair of functions $(x(s), \delta(s))$ satisfying $F(x(s); \delta(s)) = 0$ or, equivalently, $(x(s), \delta(s)) \in F^{-1}$. While there are infinitely many such pairs, a simple way to select a member of this family is to differentiate $F(x(s); \delta(s)) = 0$ with respect to s :

$$\frac{\partial F(x(s); \delta(s))}{\partial x} x'(s) + \frac{\partial F(x(s); \delta(s))}{\partial \delta} \delta'(s) = 0. \quad (10)$$

This differential equation in two unknowns $x'(s)$ and $\delta'(s)$ must be satisfied in order to remain “on path.” One possible approach for tracing out the path in F^{-1} is to solve equation (10) for the ratio $\frac{x'(s)}{\delta'(s)} = -\frac{\partial F(x(s); \delta(s))/\partial \delta}{\partial F(x(s); \delta(s))/\partial x}$ that indicates the direction of the next step along the path from s to $s + ds$. This approach, however, fails because the ratio switches from $+\infty$ to $-\infty$ at points such as B in Figure 1. So instead of solving for the ratio, we simply solve for each term of the ratio. This insight implies that the graph F^{-1} in Figure 1 is the solution to the following system of differential equations:

$$x'(s) = \frac{\partial F(x(s); \delta(s))}{\partial \delta}, \quad \delta'(s) = -\frac{\partial F(x(s); \delta(s))}{\partial x}. \quad (11)$$

These are the so-called basic differential equations for our example. They reduce the task of

tracing out the set of solutions to solving a system of differential equations. Given an initial condition this can be done with a variety of methods (see Chapter 10 of Judd 1998). If $\delta = 0$, then $F(x; \delta) = 0$ is easily solved for $x = 0.5$, thereby providing an initial condition (point A in Figure 1). From there the homotopy algorithm uses the basic differential equations to determine the next step along the path. It continues to follow—step-by-step—the path until it reaches $\delta = 1$ (point D). In our example, the auxiliary variable s is decreasing from point A to point D . Therefore, whenever $\delta'(s)$ switches sign from negative to positive (point B), the path is bending backward and there are multiple solutions. Conversely, whenever the sign of $\delta'(s)$ switches back from positive to negative (point C), the path is bending forward.

Returning to our model of learning and forgetting, let $\mathbf{x} = (\mathbf{V}^*, \mathbf{p}^*)$ denote the $2M^2$ endogenous variables. Our goal is to explore $\mathbf{F}^{-1}(\rho)$, a slice of the graph of the equilibrium correspondence. Proceeding as in our example, a parametric path is a set of functions $(\mathbf{x}(s), \delta(s)) \in \mathbf{F}^{-1}(\rho)$. Differentiating $\mathbf{F}(\mathbf{x}(s); \delta(s), \rho) = \mathbf{0}$ with respect to s yields the necessary conditions for remaining on path:

$$\frac{\partial \mathbf{F}(\mathbf{x}(s); \delta(s), \rho)}{\partial \mathbf{x}} \mathbf{x}'(s) + \frac{\partial \mathbf{F}(\mathbf{x}(s); \delta(s), \rho)}{\partial \delta} \delta'(s) = \mathbf{0}, \quad (12)$$

where $\frac{\partial \mathbf{F}(\mathbf{x}(s); \delta(s), \rho)}{\partial \mathbf{x}}$ is the $(2M^2 \times 2M^2)$ Jacobian, $\mathbf{x}'(s)$ and $\frac{\partial \mathbf{F}(\mathbf{x}(s); \delta(s), \rho)}{\partial \delta}$ are $(2M^2 \times 1)$ vectors, and $\delta'(s)$ is a scalar. This system of $2M^2$ differential equations in the $2M^2 + 1$ unknowns $x'_i(s)$, $i = 1, \dots, 2M^2$, and $\delta'(s)$ has a solution that obeys the basic differential equations

$$y'_i(s) = (-1)^{i+1} \det \left(\left(\frac{\partial \mathbf{F}(\mathbf{y}(s); \rho)}{\partial \mathbf{y}} \right)_{-i} \right), \quad i = 1, \dots, 2M^2 + 1, \quad (13)$$

where $\mathbf{y}(s) = (\mathbf{x}(s), \delta(s))$ and the notation $(\cdot)_{-i}$ is used to indicate that the i th column is removed from the $(2M^2 \times 2M^2 + 1)$ Jacobian $\frac{\partial \mathbf{F}(\mathbf{y}(s); \rho)}{\partial \mathbf{y}}$. Note that equation (13) reduces to equation (11) if \mathbf{x} is a scalar instead of a vector. Garcia & Zangwill (1979) and Chapter 2 of Zangwill & Garcia (1981) prove that the basic differential equations satisfy the conditions in equation (12).

The homotopy method requires that $\mathbf{F}(\mathbf{y}; \rho)$ is continuously differentiable with respect to \mathbf{y} and that the Jacobian $\frac{\partial \mathbf{F}(\mathbf{y}; \rho)}{\partial \mathbf{y}}$ has full rank at all points in $\mathbf{F}^{-1}(\rho)$. To appreciate the importance of regularity, note that if the Jacobian $\frac{\partial \mathbf{F}(\mathbf{y}(s); \rho)}{\partial \mathbf{y}}$ has less than full rank at some point $\mathbf{y}(s)$, then the determinants of all its $(2M^2 \times 2M^2)$ submatrices are zero. Hence, according to the basic differential equations (13), $y'_i(s) = 0$ for $i = 1, \dots, 2M^2 + 1$, and the algorithm stalls. On the other hand, with regularity in place, the implicit function theorem ensures that $\mathbf{F}^{-1}(\rho)$ consists only of continuous paths; paths that suddenly terminate,

endless spirals, branch points, isolated equilibria, and continua of equilibria are ruled out (see Chapter 1 of Zangwill & Garcia 1981).

While our assumed functional forms ensure continuous differentiability, we have been unable to establish regularity analytically. Indeed, we have numerical evidence suggesting that regularity can fail. In practice, failures of regularity are not a problem as long as they are confined to isolated points. Because our algorithm computes just a finite number of points along the path, it is extremely unlikely to hit an irregular point.¹³ We refer the reader to Borkovsky, Doraszelski & Kryukov (2008) for a fuller discussion of this issue and a step-by-step guide to solving dynamic stochastic games using the homotopy method.

As Result 1 in Section 4 shows, we have always been able to trace out a path in $\mathbf{F}^{-1}(\rho)$ that starts at the equilibrium for $\delta = 0$ and ends at the equilibrium for $\delta = 1$. Whenever this “main path” folds back on itself, the homotopy algorithm automatically identifies multiple equilibria. This makes it well-suited for models like ours that have multiple equilibria.

Nevertheless, the homotopy algorithm cannot be guaranteed to find all equilibria.¹⁴ The slice $\mathbf{F}^{-1}(\rho)$ may contain additional equilibria that are off the main path. These equilibria form one or more loops (see Result 1 in Section 4). We have two intuitively appealing but potentially fallible ways to try and identify additional equilibria. First, we use a large number of restarts of the P-M algorithm, often trying to “propagate” equilibria from “nearby” parameterizations. Second, and more systematically, just as we can choose δ as the homotopy parameter while keeping ρ fixed, we can choose ρ while keeping δ fixed. This allows us to “crisscross” the parameter space in an orderly fashion by using the equilibria on the δ -slices as initial conditions to generate ρ -slices. A ρ -slice must either intersect with all δ -slices or lead us to an additional equilibrium that, in turn, gives us an initial condition to generate an additional δ -slice. We continue this process until all the ρ - and δ -slices “match up” (for details see Grieco 2008).

3.2 Pakes & McGuire (1994) algorithm

While the homotopy method has advantages in finding multiple equilibria, it cannot stand alone. The P-M algorithm (or some other means for solving a system of nonlinear equations) is necessary to compute a starting point for our homotopy algorithm.

Recall that $V_2(\mathbf{e}) = V_1(\mathbf{e}^{[2]})$ and $p_2(\mathbf{e}) = p_1(\mathbf{e}^{[2]})$ for each state \mathbf{e} in a symmetric equilibrium and it therefore suffices to determine \mathbf{V} and \mathbf{p} , the value and policy functions of firm 1. The P-M algorithm is iterative. An iteration cycles through the states in some predetermined order and updates \mathbf{V} and \mathbf{p} as it progresses from one iteration to the next.

¹³Our programs use Hompack (Watson, Billups & Morgan 1987, Watson, Sosonkina, Melville, Morgan & Walker 1997) written in Fortran 90. They are available from the authors upon request.

¹⁴Unless the system of equations that defines them happens to be polynomial; see Judd & Schmedders (2004) for some early efforts along this line.

The strategic situation firms face in setting prices in state \mathbf{e} is similar to a static game if the value of continued play is taken as given. The P-M algorithm computes the best reply of firm 1 against $p(\mathbf{e}^{[2]})$ in this game and uses it to update the value and policy functions of firm 1 in state \mathbf{e} .

More formally, let $h_1(\mathbf{e}, p_1, p(\mathbf{e}^{[2]}), \mathbf{V})$ be the maximand in the Bellman equation (1) after symmetry is imposed. The best reply of firm 1 against $p(\mathbf{e}^{[2]})$ in state \mathbf{e} is given by

$$G_{\mathbf{e}}^2(\mathbf{V}, \mathbf{p}) = \arg \max_{p_1} h_1(\mathbf{e}, p_1, p(\mathbf{e}^{[2]}), \mathbf{V}) \quad (14)$$

and the value associated with it is

$$G_{\mathbf{e}}^1(\mathbf{V}, \mathbf{p}) = \max_{p_1} h_1(\mathbf{e}, p_1, p(\mathbf{e}^{[2]}), \mathbf{V}). \quad (15)$$

Write the collection of equations (14) and (15) for all states $\mathbf{e} \in \{1, \dots, M\}^2$ as

$$\mathbf{G}(\mathbf{V}, \mathbf{p}) = \begin{pmatrix} G_{(1,1)}^1(\mathbf{V}, \mathbf{p}) \\ G_{(2,1)}^1(\mathbf{V}, \mathbf{p}) \\ \vdots \\ G_{(M,M)}^2(\mathbf{V}, \mathbf{p}) \end{pmatrix}. \quad (16)$$

Given an initial guess $\mathbf{x}^0 = (\mathbf{V}^0, \mathbf{p}^0)$, the P-M algorithm executes the iteration

$$\mathbf{x}^{k+1} = \mathbf{G}(\mathbf{x}^k), \quad k = 0, 1, 2, \dots,$$

until the changes in the value and policy functions of firm 1 are deemed small (or a failure to converge is diagnosed).

The P-M algorithm does not lend itself to computing multiple equilibria. To identify more than one equilibrium (for a given parameterization of the model), it must be restarted from different initial guesses. But different initial guesses may or may not lead to different equilibria. This, however, still understates the severity of the problem. Whenever our dynamic stochastic game has multiple equilibria, the P-M algorithm cannot compute a substantial fraction of them even if an arbitrarily large number of initial guesses are tried.

The problem is this. The P-M algorithm continues to iterate until it reaches a fixed point $\mathbf{x} = \mathbf{G}(\mathbf{x})$. A necessary condition for convergence is local stability of the fixed point. Consider the $(2M^2 \times 2M^2)$ Jacobian $\frac{\partial \mathbf{G}(\mathbf{x})}{\partial \mathbf{x}}$ at the fixed point and let $\varrho\left(\frac{\partial \mathbf{G}(\mathbf{x})}{\partial \mathbf{x}}\right)$ be its spectral radius.¹⁵ The fixed point is locally stable under the P-M algorithm if $\varrho\left(\frac{\partial \mathbf{G}(\mathbf{x})}{\partial \mathbf{x}}\right) < 1$, i.e., if

¹⁵Let \mathbf{A} be an arbitrary matrix and $\varsigma(\mathbf{A})$ the set of its eigenvalues. The spectral radius of \mathbf{A} is $\varrho(\mathbf{A}) = \max\{|\lambda| : \lambda \in \varsigma(\mathbf{A})\}$.

all eigenvalues are within the complex unit circle. Given local stability, the P-M algorithm converges provided that the initial guess is close (perhaps very close) to the fixed point. Conversely, if $\varrho\left(\frac{\partial\mathbf{G}(\mathbf{x})}{\partial\mathbf{x}}\right) \geq 1$, then the fixed point is unstable and the P-M algorithm cannot compute it. The following proposition identifies a subset of equilibria that the P-M algorithm cannot compute.

Proposition 1 *Let $(\mathbf{x}(s), \delta(s)) \in \mathbf{F}^{-1}(\rho)$ be a parametric path of equilibria. (i) If $\delta'(s) \leq 0$, then $\varrho\left(\frac{\partial\mathbf{G}(\mathbf{x}(s))}{\partial\mathbf{x}}\Big|_{(\delta(s), \rho)}\right) \geq 1$ and the equilibrium $\mathbf{x}(s)$ is unstable under the P-M algorithm. (ii) Moreover, the equilibrium $\mathbf{x}(s)$ remains unstable with either dampening or extrapolation applied to the P-M algorithm.*

Part (i) of Proposition 1 establishes that the P-M algorithm cannot compute equilibria on any part of the path for which $\delta'(s) \leq 0$. Whenever $\delta'(s)$ switches sign from positive to negative, the main path connecting the equilibrium at $\delta = 0$ with the equilibrium at $\delta = 1$ bends backward and multiple equilibria arise. Conversely, whenever the sign of $\delta'(s)$ switches back from negative to positive, the main path bends forward. Hence, for a fixed forgetting rate $\delta(s)$, in between two equilibria for which $\delta'(s) > 0$ lies a third—necessarily unstable—equilibrium for which $\delta'(s) \leq 0$. Similarly, a loop has equilibria for which $\delta'(s) > 0$ and equilibria for which $\delta'(s) \leq 0$. Consequently, if we have multiple equilibria (for a given parameterization of the model), then the P-M algorithm can compute at best 1/2 to 2/3 of them.

Dampening and extrapolation are often applied to the P-M algorithm in the hope of improving its likelihood or speed of convergence. The iteration

$$\mathbf{x}^{k+1} = \omega\mathbf{G}(\mathbf{x}^k) + (1 - \omega)\mathbf{x}^k, \quad k = 0, 1, 2, \dots,$$

is said to be dampened if $\omega \in (0, 1)$ and extrapolated if $\omega \in (1, \infty)$. Part (ii) of Proposition 1 establishes the futility of these attempts.¹⁶

The ability of the P-M algorithm to provide a reasonably complete picture of the set of solutions to the model is limited beyond the scope of Proposition 1. Our numerical analysis indicates that the P-M algorithm cannot compute equilibria on some part of the path for which $\delta'(s) > 0$:

Proposition 2 *Let $(\mathbf{x}(s), \delta(s)) \in \mathbf{F}^{-1}(\rho)$ be a parametric path of equilibria. Even if $\delta'(s) > 0$ we may have $\varrho\left(\frac{\partial\mathbf{G}(\mathbf{x}(s))}{\partial\mathbf{x}}\Big|_{(\delta(s), \rho)}\right) \geq 1$ so that the equilibrium $\mathbf{x}(s)$ is unstable under the P-M algorithm.*

¹⁶Dampening and extrapolation may, of course, still be helpful in computing equilibria for which $\delta'(s) > 0$.

In the Online Appendix we prove Proposition 2 by way of an example and illustrate equilibria of our model that the P-M algorithm cannot compute.

As is well-known, not all Nash equilibria of static games are stable under best reply dynamics (see Chapter 1 of Fudenberg & Tirole 1991).¹⁷ Since the P-M algorithm incorporates best reply dynamics, it is reasonable to expect that this limits its usefulness. In the Online Appendix we argue that this is not the case. More precisely, we show that, holding fixed the value of continued play, the best reply dynamics are contractive and therefore converge to a unique fixed point irrespective of the initial guess. The value function iteration also is contractive holding fixed the policy function. Hence, each of the two building blocks of the P-M algorithm “works.” What makes it impossible to obtain a substantial fraction of equilibria is the interaction of value function iteration with best reply dynamics.

The P-M algorithm is a pre-Gauss-Jacobi method. The subsequent literature has instead sometimes used a pre-Gauss-Seidel method (Benkard 2004, Doraszelski & Judd 2004). Whereas a Gauss-Jacobi method replaces the old guesses for the value and policy functions with the new guesses at the end of an iteration after all states have been visited, a Gauss-Seidel method updates after each state. This has the advantage that “information” is used as soon as it becomes available (see Chapters 3 and 5 of Judd 1998). We have been unable to prove that Proposition 1 carries over to this alternative algorithm. We note, however, that the Stein-Rosenberg theorem (see Proposition 6.9 in Section 2.6 of Bertsekas & Tsitsiklis 1997) asserts that for certain systems of linear equations, if the Gauss-Jacobi algorithm fails to converge, then so does the Gauss-Seidel algorithm. Hence, it does not seem reasonable to presume that the Gauss-Seidel variant of the P-M algorithm is immune to the difficulties the original algorithm suffers.

4 Equilibrium correspondence

This section provides an overview of the equilibrium correspondence. In the absence of organizational forgetting, C-R establish uniqueness of equilibrium. Theorem 2.2 in C-R extends to our model:

Proposition 3 *If organizational forgetting is either absent ($\delta = 0$) or certain ($\delta = 1$), then there is a unique equilibrium.*¹⁸

The cases of $\delta = 0$ and $\delta = 1$ are special in that they ensure that movements through the state space are unidirectional. When $\delta = 0$, a firm can never move “backward” to a lower

¹⁷More generally, in static games, Nash equilibria of degree -1 are unstable under any Nash dynamics, i.e., dynamics with rest points that coincide with Nash equilibria, including replicator and smooth fictitious play dynamics (Demichelis & Germano 2002).

¹⁸Proposition 3 pertains to both symmetric and asymmetric equilibria.

state and when $\delta = 1$, it can never move “forward” to a higher state. Hence, backward induction can be used to establish uniqueness of equilibrium (see Section 7 for details). In contrast, when $\delta \in (0, 1)$, a firm can move in either direction. These bidirectional movements break the backward induction and make multiple equilibria possible:

Proposition 4 *If organizational forgetting is neither absent ($\delta = 0$) nor certain ($\delta = 1$), then there may be multiple equilibria.*

Figure 2 proves the proposition and illustrates the extent of multiplicity. It shows the number of equilibria that we have identified for each combination of progress ratio ρ and forgetting rate δ . Darker shades indicate more equilibria. As can be seen, we have found up to nine equilibria for some values of ρ and δ . Multiplicity is especially pervasive for forgetting rates δ in the empirically relevant range below 0.1.

In dynamic stochastic games with finite actions, Herings & Peeters (2004) have shown that generically the number of MPE is odd. While they consider both symmetric and asymmetric equilibria, in a two-player game with symmetric primitives such as ours, asymmetric equilibria occur in pairs. Hence, their result immediately implies that generically the number of symmetric equilibria is odd in games with finite actions. Figure 2 suggests that this carries over to our setting with continuous actions.

In order to understand the geometry of how multiple equilibria arise we take a close look at the slices of the graph of the equilibrium correspondence that our homotopy algorithm computes.

Result 1 *The slice $\mathbf{F}^{-1}(\rho)$ contains a unique path that connects the equilibrium at $\delta = 0$ with the equilibrium at $\delta = 1$. In addition, $\mathbf{F}^{-1}(\rho)$ may contain (one or more) loops that are disjoint from this “main path” and from each other.*

Figure 3 illustrates Result 1. To explain this figure, recall that an equilibrium consists of a value function \mathbf{V}^* and a policy function \mathbf{p}^* and is thus an element of a high-dimensional space. To succinctly represent it, we proceed in two steps. First, we use \mathbf{p}^* to construct the probability distribution over next period’s state \mathbf{e}' given this period’s state \mathbf{e} , i.e., the transition matrix that characterizes the Markov process of industry dynamics. We compute the transient distribution over states in period t , $\boldsymbol{\mu}^t(\cdot)$, starting from state $(1, 1)$ in period 0. This tells us how likely each possible industry structure is in period t , given that both firms began at the top of their learning curves. In addition, we compute the limiting (or ergodic) distribution over states, $\boldsymbol{\mu}^\infty(\cdot)$.¹⁹ The transient distributions capture short-run dynamics and the limiting distribution captures long-run (or steady-state) dynamics.

¹⁹Let \mathbf{P} be the $M^2 \times M^2$ transition matrix. The transient distribution in period t is given by $\boldsymbol{\mu}^t = \boldsymbol{\mu}^0 \mathbf{P}^t$, where $\boldsymbol{\mu}^0$ is the $1 \times M^2$ initial distribution and \mathbf{P}^t the t^{th} matrix power of \mathbf{P} . If $\delta \in (0, 1)$, then the Markov process is irreducible because logit demand implies that the probability moving forward is always nonzero.

Second, we use the transient distribution over states in period t , $\boldsymbol{\mu}^t(\cdot)$, to compute the expected Herfindahl index

$$H^t = \sum_{\mathbf{e}} (D_1^*(\mathbf{e})^2 + D_2^*(\mathbf{e})^2) \boldsymbol{\mu}^t(\mathbf{e}).$$

The time path of H^t summarizes the implications of learning and forgetting for industry dynamics. If the industry evolves asymmetrically, then $H^t > 0.5$. The maximum expected Herfindahl index

$$H^\wedge = \max_{t \in \{1, \dots, 100\}} H^t$$

is a summary measure of short-run industry concentration. The limiting expected Herfindahl index H^∞ , computed using $\boldsymbol{\mu}^\infty(\cdot)$ instead of $\boldsymbol{\mu}^t(\cdot)$, is a summary measure of long-run industry concentration. If $H^\infty > 0.5$, then an asymmetric industry structure persists.

In Figure 3 we visualize $\mathbf{F}^{-1}(\rho)$ for a variety of progress ratios by plotting H^\wedge (dashed line) and H^∞ (solid line). As can be seen, multiple equilibria arise whenever the main path folds back on itself. Moreover, there is one loop for $\rho \in \{0.75, 0.65, 0.55, 0.15, 0.05\}$, two loops for $\rho \in \{0.85, 0.35\}$, and three loops for $\rho = 0.95$, thus adding further multiplicity.

Figure 3 is not necessarily a complete picture of the equilibria to our model. As discussed in Section 3.1, no algorithm is guaranteed to find all equilibria. We do find all equilibria along the main path and we have been successful in finding a number of loops. But other loops may exist because, in order to trace out a loop, we must somehow compute at least one equilibrium on the loop, and doing so is problematic.

Types of equilibria. Despite the multiplicity, the equilibria of our game exhibit the four typical patterns shown in Figure 4.²⁰ The parameter values are $\rho = 0.85$ and $\delta \in \{0, 0.0275, 0.08\}$; they represent the median progress ratio across a wide array of empirical studies combined with the cases of no, low, and high organizational forgetting. One should recognize that the typical patterns, helpful as they are in understanding the range of behaviors that can occur, lie on a continuum and thus morph into each other in complicated ways as we change the parameter values.

The upper left panel of Figure 4 is typical for what we call a *flat equilibrium without well* ($\rho = 0.85$, $\delta = 0$). The policy function is very even over the entire state space. In particular, the price that a firm charges in equilibrium is fairly insensitive to its rival's stock of know-how. The upper right panel shows a *flat equilibrium with well* ($\rho = 0.85$,

That is, all its states belong to a single closed communicating class and the $1 \times M^2$ limiting distribution $\boldsymbol{\mu}^\infty$ solves the system of linear equations $\boldsymbol{\mu}^\infty = \boldsymbol{\mu}^\infty \mathbf{P}$. If $\delta = 0$ ($\delta = 1$), then there is also a single closed communicating class, but its sole member is state (M, M) $((1, 1))$.

²⁰The value functions corresponding to the policy functions in Figure 4 can be found in the Online Appendix where we also provide tables of the value and policy functions for ease of reference.

$\delta = 0.0275$). While the policy function remains even over most of the state space, price competition is intense during the industry’s birth. This manifests itself as a “well” in the neighborhood of state $(1, 1)$.

The lower left panel of Figure 4 exemplifies a *trenchy equilibrium* ($\rho = 0.85, \delta = 0.0275$). The parameter values are the same as for the flat equilibrium with well, thereby providing an instance of multiplicity.²¹ The policy function is uneven and exhibits a “trench” along the diagonal of the state space. This trench starts in state $(1, 1)$ and extends beyond the bottom of the learning curve in state (m, m) all the way to state (M, M) . Hence, in a trenchy equilibrium, price competition between firms with similar stocks of know-how is intense but abates once firms become asymmetric. Finally, the lower right panel illustrates an *extra-trenchy equilibrium* ($\rho = 0.85, \delta = 0.08$). The policy function has not only a diagonal trench but also trenches parallel to the edges of the state space. In these sideways trenches the leader competes aggressively with the follower.

Sunspots. For a progress ratio of $\rho = 1$ the marginal cost of production is constant at $c(1) = \dots = c(M) = \kappa$, and there are no gains from learning-by-doing. It clearly is an equilibrium for firms to disregard their stocks of know-how and set the same prices as in the Nash equilibrium of a static price-setting game (obtained by setting $\beta = 0$). Since firms’ marginal costs are constant, so are the static Nash equilibrium prices. Thus, we have an extreme example of a flat equilibrium with $p^*(\mathbf{e}) = \kappa + 2\sigma = 12$ and $V^*(\mathbf{e}) = \frac{\sigma}{1-\beta} = 21$ for all states $\mathbf{e} \in \{1, \dots, M\}^2$.

Figure 2 shows that, in case of $\rho = 1$, there are two more equilibria for a range of forgetting rates δ below 0.1. Since the state of the industry has no bearing on the primitives, we refer to these equilibria as sunspots. One of the sunspots is a trenchy equilibrium while the other one is, depending on δ , either a flat or a trenchy equilibrium. In the trenchy equilibrium the industry evolves towards an asymmetric structure where the leader charges a lower price than the follower and enjoys a higher probability of making a sale. Consequently, the net present value of cash flows to the leader exceeds that to the follower. The value in state $(1, 1)$, however, is lower than in the static Nash equilibrium, i.e., $V^*(1, 1) < 21$.²² This indicates that value is destroyed as firms fight for dominance.

The existence of sunspots and the fact that these equilibria persist for $\rho \approx 1$ suggests that the concept of MPE is richer than one may have thought. Besides describing the

²¹As can be seen in the upper right panel of Figure 3, the main path in $\mathbf{F}^{-1}(0.85)$ bends back on itself at $\delta = 0.0275$, and there are three equilibria for slightly lower values of δ and only one for slightly higher values. This particular parameterization (if not the pattern of behavior it generates) is therefore almost nongeneric in that it approximates the isolated occurrence of an even number of equilibria. Due to the limited precision of our homotopy algorithm, we have indeed been unable to find a third equilibrium.

²²For example, if $\delta = 0.0275$, then $V^*(28, 21) = 25.43$ and $p^*(28, 21) = 12.33$ for the leader, $V^*(21, 28) = 22.39$ and $p^*(21, 28) = 12.51$ for the follower, and $V^*(1, 1) = 19.36$.

physical environment of the industry the state serves as a summary of the history of play: A larger stock of know-how indicates that—on average—a firm has won more sales than its rival, with the likely reason being that the firm has charged lower prices. Hence, by conditioning their current behavior on the state, firms implicitly condition on the history of play. The difference with a subgame perfect equilibrium is that there firms have the entire history of play at their disposal whereas here they have but a crude indication of it. Nevertheless, “barely” payoff-relevant state variables (such as firms’ stocks of know-how if $\rho \approx 1$) open the door for bootstrap-type equilibria as familiar from repeated games to arise in Markov-perfect settings.

In sum, accounting for organizational forgetting in a model of learning-by-doing leads to multiple equilibria and a rich array of pricing behaviors. In the next section we explore what these behaviors entail for industry dynamics, both in the short run and in the long run.

5 Industry dynamics

Figures 5 and 6 display the transient distribution in period 8 and 32, respectively, and Figure 7 displays the limiting distribution for our four typical cases.²³ In the flat equilibrium without well ($\rho = 0.85$, $\delta = 0$, upper left panels), the transient and limiting distributions are unimodal. The most likely industry structure is symmetric. For example, the modal state is (5, 5) in period 8, (9, 9) in period 16, (17, 17) in period 32, and (30, 30) in period 64. Turning from the short run to the long run, the industry is sure to remain in state (30, 30) because with logit demand a firm always has a positive probability of making a sale irrespective of its own price and that of its rival so that, in the absence of organizational forgetting, both firms must eventually reach the bottom of their learning curves.²⁴ In short, the industry starts symmetric and stays symmetric.

By contrast, in the flat equilibrium with well ($\rho = 0.85$, $\delta = 0.0275$, upper right panels) the transient distributions are first bimodal and then unimodal as is the limiting distribution. The modal states are (1, 8) and (8, 1) in period 8, (4, 11) and (11, 4) in period 16, (9, 14) and (14, 9) in period 32, but the modal state is (17, 17) in period 64 and the modal states of the limiting distribution are (24, 25) and (25, 24). Thus, as time passes, firms end up competing on equal footing. In sum, the industry evolves first towards an asymmetric structure and then towards a symmetric structure. As we discuss in detail in the Section 6, the well serves to build, but not to defend, a competitive advantage.

²³To avoid clutter, we do not graph states that have probability of less than 10^{-4} .

²⁴The absence of persistent asymmetries is not an artifact of our functional forms. C-R point out that it holds true as long as the support of demand is unbounded (see their Assumption 1(a) and footnote 6 on p. 1118).

While the modes of the transient distributions are more separated and pronounced in the trenchy equilibrium ($\rho = 0.85$, $\delta = 0.0275$, lower left panels) than in the flat equilibrium with well, the dynamics of the industry are similar at first. Unlike in the flat equilibrium with well, however, the industry continues to evolve towards an asymmetric structure. The modal states are (14, 21) and (21, 14) in period 64 and (21, 28) and (28, 21) in the limiting distribution. Although the follower reaches the bottom of its learning curve in the long run and attains cost parity with the leader, asymmetries persist because the diagonal trench serves to build and to defend a competitive advantage.

In the extra-trenchy equilibrium ($\rho = 0.85$, $\delta = 0.08$, lower right panels) the sideways trench renders it unlikely that the follower ever makes it down from the top of its learning curve. The transient and limiting distributions are bimodal, and the most likely industry structure is extremely asymmetric. The modal states are (1, 7) and (7, 1) in period 8, (1, 10) and (10, 1) in period 16, (1, 15) and (15, 1) in period 32, and (1, 19) and (19, 1) in period 64, and (1, 26) and (26, 1) in the limiting distribution. In short, one firm acquires a competitive advantage early on and maintains it with an iron hand.

Returning to Figure 3, the maximum expected Herfindahl index H^\wedge (dashed line) and the limiting expected Herfindahl index H^∞ (solid line) highlight the fundamental economics of organizational forgetting. If forgetting is sufficiently weak ($\delta \approx 0$), then asymmetries may arise but cannot persist, i.e., $H^\wedge \geq 0.5$ and $H^\infty \approx 0.5$. Moreover, if asymmetries arise in the short run, they are modest. If forgetting is sufficiently strong ($\delta \approx 1$), then asymmetries cannot arise in the first place, i.e., $H^\wedge \approx H^\infty \approx 0.5$ because forgetting stifles investment in learning altogether.²⁵ But, for intermediate degrees of forgetting, asymmetries arise and persist. These asymmetries can be so pronounced that the leader is virtually a monopolist.

Since the Markov process of industry dynamics is irreducible for $\delta \in (0, 1)$, the follower must eventually overtake the leader. The limiting expected Herfindahl index H^∞ may be a misleading measure of long-run industry concentration if such leadership reversals happen frequently. This, however, is not the case: Leadership reversals take a long time to occur when H^∞ is high. To establish this, define $\tau(e_1, e_2)$ to be the first-passage time into the set $\{(\tilde{e}_1, \tilde{e}_2) | \tilde{e}_1 \leq \tilde{e}_2\}$ if $e_1 \geq e_2$ or $\{(\tilde{e}_1, \tilde{e}_2) | \tilde{e}_1 \geq \tilde{e}_2\}$ if $e_1 \leq e_2$. That is, $\tau(\mathbf{e})$ is the expected time it takes the industry to move from state \mathbf{e} below (or on) the diagonal of the state space, where firm 1 leads and firm 2 follows, to state $\tilde{\mathbf{e}}$ above (or on) it, where firm 1 follows and firm 2 leads. Taking the average with respect to the limiting distribution yields the summary measure

$$\tau^\infty = \sum_{\mathbf{e}} \tau(\mathbf{e}) \boldsymbol{\mu}^\infty(\mathbf{e}).$$

For the trenchy and extra-trenchy equilibria $\tau^\infty = 295$ and $\tau^\infty = 83,406$, respectively, indi-

²⁵We further document this investment stifling in the Online Appendix.

cating that a leadership reversal takes a long time to occur. Hence, the asymmetry captured by H^∞ persists. In the Online Appendix we plot the expected time to a leadership reversal τ^∞ in the same format as Figure 3. Just like H^∞ , τ^∞ is largest for intermediate degrees of forgetting. Moreover, τ^∞ is of substantial magnitude, easily reaching and exceeding 1,000 periods. Asymmetries are therefore persistent in our model because the expected time until the leader and the follower switch roles is (perhaps very) long.

We caution the reader that the absence of persistent asymmetries for small forgetting rates δ in Figure 3 may be an artifact of the finite size of the state space ($M = 30$ in our baseline parameterization). Given $\delta = 0.01$, say, $\Delta(30) = 0.26$ and organizational forgetting is so weak that the industry is sure to remain in or near state $(30, 30)$. This eliminates bidirectional movements sufficiently completely so as to restore the backward induction logic that underlies uniqueness of equilibrium for the extreme case of $\delta = 0$ (see Proposition 3). We show in the Online Appendix that, increasing M , whilst holding fixed δ , facilitates persistent asymmetries as the industry becomes more likely to remain in the interior of the state space. Furthermore, as emphasized in Section 1, shut-out model elements can give rise to persistent asymmetries even in the absence of organizational forgetting. We explore this issue further in Section 8.

To summarize, contrary to what one might expect, organizational forgetting does not negate learning-by-doing. Rather, as can be seen in Figure 3, over a range of progress ratios ρ above 0.6 and forgetting rates δ below 0.1, learning and forgetting reinforce each other. Starting from the absence of both learning ($\rho = 1$) and forgetting ($\delta = 0$), a steeper learning curve (lower progress ratio) tends to give rise to a more asymmetric industry structure just as a higher forgetting rate does. In the next section we analyze the pricing behavior that drives these dynamics.

6 Pricing behavior

Re-writing equation (5) shows that firm 1's price in state \mathbf{e} satisfies

$$p^*(\mathbf{e}) = c^*(\mathbf{e}) + \frac{\sigma}{1 - D_1^*(\mathbf{e})}, \quad (17)$$

where the *virtual marginal cost*,

$$c^*(\mathbf{e}) = c(e_1) - \beta\phi^*(\mathbf{e}), \quad (18)$$

equals the *actual marginal cost*, $c(e_1)$, minus the *discounted prize*, $\beta\phi^*(\mathbf{e})$, from winning the current period's sale. The prize, in turn, is the difference in the value of continued play to

firm 1 if it wins the sale, $\bar{V}_1^*(\mathbf{e})$, versus if it loses the sale, $\bar{V}_2^*(\mathbf{e})$:

$$\phi^*(\mathbf{e}) = \bar{V}_1^*(\mathbf{e}) - \bar{V}_2^*(\mathbf{e}). \quad (19)$$

Note that, irrespective of the forgetting rate δ , the equilibrium of our dynamic stochastic game reduces to the static Nash equilibrium if firms are myopic. Setting $\beta = 0$ in equations (17) and (18) gives the usual FOC for a static price-setting game with logit demand:

$$p^\dagger(\mathbf{e}) = c(e_1) + \frac{\sigma}{1 - D_1^\dagger(\mathbf{e})}, \quad (20)$$

where $D_k^\dagger(\mathbf{e}) = D_k(p^\dagger(\mathbf{e}), p^\dagger(\mathbf{e}^{[2]}))$ denotes the probability that, in the static Nash equilibrium, the buyer purchases from firm $k \in \{1, 2\}$ in state \mathbf{e} . Thus, if $\beta = 0$, then $p^*(\mathbf{e}) = p^\dagger(\mathbf{e})$ and $V^*(\mathbf{e}) = D_1^\dagger(\mathbf{e}) (p^\dagger(\mathbf{e}) - c(e_1))$ for all states $\mathbf{e} \in \{1, \dots, M\}^2$.

6.1 Price bounds

Comparing equations (17) and (20) shows that equilibrium prices $p^*(\mathbf{e})$ and $p^*(\mathbf{e}^{[2]})$ coincide with the prices that obtain in a static Nash equilibrium with costs equal to virtual marginal costs $c^*(\mathbf{e})$ and $c^*(\mathbf{e}^{[2]})$. Static Nash equilibrium prices are increasing in either firm's cost (Vives 1999, p. 35). Therefore, if both firms' prizes are nonnegative, static Nash equilibrium prices are an upper bound on equilibrium prices, i.e., if $\phi^*(\mathbf{e}) \geq 0$ and $\phi^*(\mathbf{e}^{[2]}) \geq 0$, then $p^*(\mathbf{e}) \leq p^\dagger(\mathbf{e})$ and $p^*(\mathbf{e}^{[2]}) \leq p^\dagger(\mathbf{e}^{[2]})$.

A sufficient condition for $\phi^*(\mathbf{e}) \geq 0$ for each state \mathbf{e} is that the value function $V^*(\mathbf{e})$ is nondecreasing in e_1 and nonincreasing in e_2 . Intuitively, it should not hurt firm 1 if it moves down its learning curve and it should not benefit firm 1 if firm 2 moves down its learning curve. While neither we nor C-R have succeeded in proving it, our computations show that this intuition is valid in the absence of organizational forgetting:

Result 2 *If organizational forgetting is absent ($\delta = 0$), then $p^*(\mathbf{e}) \leq p^\dagger(\mathbf{e})$ for all $\mathbf{e} \in \{1, \dots, M\}^2$.*

Result 2 highlights the fundamental economics of learning-by-doing: As long as improvements in competitive position are valuable, firms use price cuts as investments to achieve them.

We complement Result 2 by establishing a lower bound on equilibrium prices in states where at least one of the two firms has reached the bottom of its learning curve:

Proposition 5 *If organizational forgetting is absent ($\delta = 0$), then (i) $p^*(\mathbf{e}) = p^\dagger(\mathbf{e}) = p^\dagger(m, m) > c(m)$ for all $\mathbf{e} \in \{m, \dots, M\}^2$ and (ii) $p^*(\mathbf{e}) > c(m)$ for all $e_1 \in \{m, \dots, M\}$ and $e_2 \in \{1, \dots, m - 1\}$.*

Part (i) of Proposition 5 sharpens Theorem 4.3 in C-R by showing that once both firms have reached the bottom of their learning curves, equilibrium prices revert to static Nash levels. To see why, note that, given $\delta = 0$, the prize reduces to $\phi^*(\mathbf{e}) = V^*(e_1 + 1, e_2) - V^*(e_1, e_2 + 1)$. But beyond the bottom of their learning curves, firms' competitive positions can neither improve nor deteriorate. Hence, as we show in the proof of the proposition, $V^*(\mathbf{e}) = V^*(\mathbf{e}')$ for all $\mathbf{e}, \mathbf{e}' \in \{m, \dots, M\}^2$, so that the advantage-building and advantage-defending motives disappear, the prize is zero, and equilibrium prices revert to static Nash levels. This rules out trenches penetrating into this region of the state space. Trenchy and extra-trenchy equilibria therefore cannot arise in the absence of organizational forgetting.

Part (ii) of Proposition 5 restates Theorem 4.3 in C-R for the situation where the leader but not the follower has reached the bottom of its learning curve. The leader no longer has an advantage-building motive but continues to have an advantage-defending motive. This raises the possibility that the leader uses price cuts to delay the follower in moving down its learning curve. The proposition shows that there is a limit to how aggressively the leader defends its advantage: below-cost pricing is never optimal in the absence of organizational forgetting.

The story changes dramatically in the presence of organizational forgetting. The equilibrium may exhibit soft competition in some states and price wars in other states. Consider the trenchy equilibrium ($\rho = 0.85$, $\delta = 0.0275$). The upper bound in Result 2 fails in state $(22, 20)$ where the leader charges 6.44 and the follower charges 7.60, significantly above its static Nash equilibrium price of 7.30. The follower's high price stems from its prize of -1.04 . This prize, in turn, reflects that, if the follower wins the sale, then the industry most likely moves to state $(22, 21)$ and thus closer to the brutal price competition on the diagonal of the state space. Indeed, the follower's value function decreases from 20.09 in state $(22, 20)$ to 19.56 in state $(22, 21)$. To avoid this undesirable possibility, the follower charges a high price. The lower bound in Proposition 5 fails in state $(20, 20)$ where both firms charge 5.24 as compared to a marginal cost of 5.30. The prize of 2.16 makes it worthwhile to price below cost even beyond the bottom of the learning curve because "in the trench" winning the current sale confers a lasting advantage.

This discussion provides the instances that prove the next two propositions.

Proposition 6 *If organizational forgetting is present ($\delta > 0$), then we may have $p^*(\mathbf{e}) > p^\dagger(\mathbf{e})$ for some $\mathbf{e} \in \{1, \dots, M\}^2$.*

Figure 8 illustrates Proposition 6 by plotting the share of equilibria that violate the upper bound in Result 2.²⁶ Darker shades indicate higher shares. As can be seen, the upper

²⁶To take into account the limited precision of our computations, we take the upper bound to be violated if $p^*(\mathbf{e}) > p^\dagger(\mathbf{e}) + \epsilon$ for some $\mathbf{e} \in \{1, \dots, M\}^2$, where ϵ is positive but small. Specifically, we set $\epsilon = 10^{-2}$, so

bound continues to hold if organizational forgetting is very weak ($\delta \approx 0$) and possibly also if learning-by-doing is very weak ($\rho \approx 1$). Apart from these extremes (and a region around $\rho = 0.45$ and $\delta = 0.25$), at least some, if not all, equilibria entail at least one state where equilibrium prices exceed static Nash equilibrium prices.

Taken alone, Proposition 6 suggests that organizational forgetting makes firms less aggressive. This makes sense: After all, why invest in improvements in competitive position when they are transitory? But organizational forgetting can also be a source of aggressive pricing behavior:

Proposition 7 *If organizational forgetting is present ($\delta > 0$), then we may have $p^*(\mathbf{e}) \leq c(m)$ for some $e_1 \in \{m, \dots, M\}$ and $e_2 \in \{1, \dots, M\}$.*

Figure 9 illustrates Proposition 7 by plotting the share of equilibria in which a firm prices below cost even though it has reached the bottom of its learning curve. Note that Figure 9 is a conservative tally of how often the lower bound in Proposition 5 fails because the lower bound in part (i) already fails if the leader charges less than its static Nash equilibrium price, not less than its marginal cost. In sum, the leader may be more aggressive in defending its advantage in the presence of organizational forgetting than in its absence. The most dramatic expression of this aggressive pricing behavior are the diagonal trenches that are the defining feature of trenchy and extra-trenchy equilibria.

6.2 Wells and trenches

This section develops intuition as to how wells and trenches can arise. Our goal is to provide insight as to whether equilibria featuring wells and trenches are economically plausible and, at least potentially, empirically relevant.

Wells. A *well*, as seen in the upper right panel of Figure 4, is a preemption battle that firms at the top of their learning curves fight. Consider our leading example of a flat equilibrium with well ($\rho = 0.85$, $\delta = 0.0275$).²⁷ Table 1 details firms' competitive positions at various points in time assuming that firm 1 leads and firm 2 follows. Having moved down the learning curve first, the leader has a lower cost and a higher prize than the follower. In the modal state (8, 1) in period 8 the leader therefore charges a lower price and enjoys a higher probability of making a sale. In time the follower also moves down the learning curve and the leader's advantage begins to erode (see modal state (11, 4) in period 16) and

that if prices are measured in dollars, then the upper bound must be violated by more than a cent. Given that the homotopy algorithm solves the system of equations up to a maximum absolute error of about 10^{-12} , Figure 8 therefore almost certainly understates the extent of violations.

²⁷As a point of comparison, we provide details on firms' competitive positions at various points in time for our leading example of a flat equilibrium without well ($\rho = 0.85$, $\delta = 0$) in the Online Appendix.

period	modal state	leader					follower				
		cost	prize	price	prob.	value	cost	prize	price	prob.	value
0	(1,1)	10.00	6.85	5.48	0.50	5.87	10.00	6.85	5.48	0.50	5.87
8	(8,1)	6.14	3.95	7.68	0.81	22.99	10.00	2.20	9.14	0.19	5.34
16	(11,4)	5.70	1.16	7.20	0.62	20.08	7.22	1.23	7.68	0.38	11.48
32	(14,9)	5.39	0.36	7.16	0.53	20.06	5.97	0.64	7.27	0.47	17.30
64	(17,17)	5.30	-0.01	7.31	0.50	20.93	5.30	-0.01	7.31	0.50	20.93
∞	(25,24)	5.30	-0.01	7.30	0.50	21.02	5.30	-0.00	7.30	0.50	21.02

Table 1: Cost, prize, price, probability of making a sale, and value. Flat equilibrium with well ($\rho = 0.85$, $\delta = 0.0275$).

state	leader					follower				
	cost	prize	price	prob.	value	cost	prize	price	prob.	value
(21, 20)	5.30	3.53	5.57	0.72	21.91	5.30	0.14	6.54	0.28	19.56
(21, 21)	5.30	2.14	5.26	0.50	19.79	5.30	2.14	5.26	0.50	19.79
(22, 20)	5.30	3.22	6.44	0.76	23.98	5.30	-1.04	7.60	0.24	20.09
(28, 21)	5.30	-0.13	7.63	0.55	25.42	5.30	-0.71	7.81	0.45	22.37
(20, 20)	5.30	2.16	5.24	0.50	19.82	5.30	2.16	5.24	0.50	19.82

Table 2: Cost, prize, price, probability of making a sale, and value. Trenchy equilibrium ($\rho = 0.85$, $\delta = 0.0275$).

eventually vanishes completely (see the modal state (17, 17) in period 64). The prizes reflect this erosion. The leader’s prize is higher than the follower’s in state (8, 1) (3.95 versus 2.20) but lower in state (11, 4) (1.16 versus 1.23). Although the leadership position is transitory, it is surely worth having. Both firms use price cuts in state (1, 1) in the hope of being the first to move down the learning curve. In the example, the prize of 6.85 justifies charging the price of 5.48 that is well below the marginal cost of 10. The well is therefore an investment in building competitive advantage.

More abstractly, a well is the outcome of an auction in state (1, 1) for the additional future profits—the prize—that accrue to the firm that makes the first sale and acquires transitory industry leadership. As equations (17)–(19) show, a firm’s price in equilibrium is virtual marginal cost marked up. Virtual marginal cost, in turn, accounts for the discounted prize from winning the current period’s sale and, getting to the essential point, the prize is the difference in the value of continued play if the firm rather than its rival wins.

Diagonal trenches. A *diagonal trench*, as seen in the lower panels of Figure 4, is a price war between symmetric or nearly symmetric firms. Extending along the entire diagonal of the state space, a diagonal trench has the curious feature that the firms compete fiercely—perhaps pricing below cost—even when they both have exhausted all gains from

learning-by-doing. Part (i) of Proposition 5 rules out this type of behavior in the absence of organizational forgetting.

Like a well, a diagonal trench serves to build a competitive advantage. Unlike a well, a diagonal trench also serves to defend a competitive advantage, thereby rendering it (almost) permanent: The follower recognizes that to seize the leadership position it would have to “cross over” the diagonal trench and struggle through another price war. Crucially this price war is a part of a MPE and, as such, a credible threat the follower cannot ignore.

The logic behind a diagonal trench has three parts. If a diagonal trench exists, then the follower does not contest the leadership position. If the follower does not contest the leadership position, then being the leader is valuable. Finally, to close the circle of logic, if being the leader is valuable, then firms price aggressively on the diagonal on the state space in a bid for the leadership position, thus giving rise to the diagonal trench. Table 2 illustrates this argument by providing details on firms’ competitive position in various states for our leading example of a trenchy equilibrium ($\rho = 0.85$, $\delta = 0.0275$).

Part 1: Trench sustains leadership. To see why the follower does not contest the leadership position consider a state such as (21, 20) where the follower has almost caught up with the leader. Suppose the follower wins the current sale. In this case the follower may leapfrog the leader if the industry moves against the odds to state (20, 21). However, the most likely possibility, with a probability of 0.32, is that the industry moves to state (21, 21). Due to the brutal price competition “in the trench” the follower’s expected cash flow in the next period decreases to $-0.02 = 0.50 \times (5.26 - 5.30)$ compared to $0.34 = 0.28 \times (6.54 - 5.30)$ if the industry had remained in state (21, 20). Suppose, in contrast, the leader wins. This completely avoids sparking a price war. Moreover, the most likely possibility, with a probability of 0.32, is that the leader enhances its competitive advantage by moving to state (22, 20). If so, the leader’s expected cash flow in the next period increases to $0.87 = 0.76 \times (6.44 - 5.30)$ compared to $0.20 = 0.72 \times (5.57 - 5.30)$ if the industry had remained in state (21, 20). Because winning the sale is more valuable to the leader than to the follower, the leader’s prize in state (21, 20) is almost 25 times larger than the follower’s and the leader underprices the follower. As a consequence, the leader defends its position with a substantial probability of 0.79. In other words, the diagonal trench sustains the leadership position.

Part 2: Leadership generates value. Because the leader underprices the follower, over time the industry moves from state (21, 20) to (or near) the modal state (28, 21) of the limiting distribution. Once there the leader underprices the follower (7.63 versus 7.81) despite cost parity and thus enjoys a higher probability of making a sale (0.55 versus 0.45). The leader’s expected cash flow in the current period is therefore $0.55 \times (7.63 - 5.30) = 1.27$ as compared to the follower’s of $0.45 \times (7.81 - 5.30) = 1.14$. Because the follower does not contest the leadership position, the leader is likely to enjoy these additional profits for a

state	leader					follower				
	cost	prize	price	prob.	value	cost	prize	price	prob.	value
(26, 1)	5.30	6.43	8.84	0.90	53.32	10.00	0.12	11.00	0.10	2.42
(26, 2)	5.30	6.21	7.48	0.88	46.31	8.50	0.21	9.44	0.12	2.55
(26, 3)	5.30	5.16	6.94	0.85	40.23	7.73	0.27	8.65	0.15	2.78
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
(26, 7)	5.30	3.04	6.14	0.73	24.76	6.34	0.58	7.15	0.27	4.16
(26, 8)	5.30	2.33	5.99	0.66	21.86	6.14	1.08	6.64	0.34	4.78
(26, 9)	5.30	1.17	6.24	0.51	19.84	5.97	1.71	6.29	0.49	6.07
(26, 10)	5.30	0.16	6.83	0.40	19.09	5.83	1.96	6.44	0.60	8.08

Table 3: Cost, prize, price, probability of making a sale, and value. Extra-trenchy equilibrium ($\rho = 0.85$, $\delta = 0.08$).

long time (recall that $\tau^\infty = 295$). Hence, being the leader is valuable.

Part 3: Value induces trench. Because being the leader is valuable, firms price aggressively on the diagonal on the state space in a bid for the leadership position. The prize is 2.16 in state (20, 20) and 2.14 in state (21, 21) and justifies charging a price of 5.24 and 5.26, respectively, even through all gains from learning-by-doing have been exhausted. This gives rise to the diagonal trench. Observe that this argument applies at every state on the diagonal because, no matter where on the diagonal the firms happen to be, winning the current sale confers a lasting advantage. The trench therefore extends along the entire diagonal of the state space.

All this can be summed up in a sentence: Building a competitive advantage creates the diagonal trench that defends the advantage and creates the prize that makes it worthwhile to fight for the leadership position. A diagonal trench is thus a self-reinforcing mechanism for gaining and maintaining market dominance.

Sideways trenches. A *sideways trench*, as seen in the lower right panel of Figure 4, is a price war between very asymmetric firms. This war is triggered when the follower starts to move down its learning curve. Table 3 provides details on firms' competitive positions in various states for our leading example of an extra-trenchy equilibrium ($\rho = 0.85$, $\delta = 0.08$). The sideways trench is evident in the decrease in the leader's price from state (26, 1) to state (26, 8) and the increase from state (26, 8) to state (26, 10). Note that the follower has little chance of making it down its learning curve as long as the probability of winning a sale is less than the probability of losing a unit of know-how through organizational forgetting. While $D_2^*(26, 1) = 0.10 > 0.08 = \Delta(1)$, we have $D_2^*(26, 2) = 0.12 < 0.15 = \Delta(2)$ and $D_2^*(26, 3) = 0.15 < 0.22 = \Delta(3)$. Hence, the leader can stall the follower at the top of its

learning curve and, indeed, the modal state of the limiting distribution is (26, 1).

The additional future profits stemming from the leader’s ability to stall the follower are the source of its large prize in state (26, 1). In state (26, 2) the prize is almost as large because by winning a sale the leader may move the industry back to state (26, 1) in the next period. The leader’s prize falls as the follower moves further down its learning curve because it takes progressively longer for the leader to force the follower back up its learning curve and because the lower cost of the follower makes it harder for the leader to do so. In the unlikely event that the follower crashes through the sideways trench in state (26, 8), the leader’s prize falls sharply. At the same time the follower’s prize rises sharply as it turns from a docile competitor into a viable threat.

A sideways trench, like a diagonal trench, is a self-reinforcing mechanism for gaining and maintaining market dominance. But whereas a diagonal trench is about fighting an imminent threat, a sideways trench is about fighting a distant threat. One can think of a sideways trench as an endogenously arising mobility barrier in the sense of Caves & Porter (1977) or the equilibrium manifestation of former Intel CEO Andy Grove’s dictum “Only the paranoid survive.”

In sum, the four types of equilibria that we have identified in Section 4 give rise to distinct yet plausible pricing behaviors and industry dynamics. Rather than impeding aggressive behavior, organizational forgetting facilitates it. In its absence, the equilibria are flat either with or without well depending on the progress ratio. Generally speaking, organizational forgetting is associated with “trenchier” equilibria, more aggressive behavior, and more concentrated industries both in the short run and in the long run.

6.3 Dominance properties

Traditional intuition suggests that learning-by-doing leads by itself to market dominance by giving a more experienced firm the ability to profitably underprice its less experienced rival. This enables the leader to widen its competitive advantage over time, thereby further enhancing its ability to profitably underprice the follower. C-R formalize this idea with “two concepts of self-reinforcing market dominance” (p. 1115): An equilibrium exhibits *increasing dominance* (ID) if $p^*(\mathbf{e}) - p^*(\mathbf{e}^{[2]}) < 0$ whenever $e_1 > e_2$ and *increasing increasing dominance* (IID) if $p^*(\mathbf{e}) - p^*(\mathbf{e}^{[2]})$ is decreasing in e_1 . If ID holds, the leader charges a lower price than the follower and therefore enjoys a higher probability of making a sale. If IID holds, the price gap between the firms widens with the length of the lead.²⁸

²⁸Athey & Schmutzler’s (2001) notion of weak increasing dominance describes the relationship between players’ states and their actions in dynamic games with deterministic state-to-state transitions and coincides with the notion of ID in C-R. Similar notions have also been used by Vickers (1986) and Budd, Harris & Vickers (1993) in dynamic investment games.

In the absence of organizational forgetting, Theorem 3.3 in C-R shows that ID and IID hold provided that the discount factor β is sufficiently close to 1 (or, alternatively, to 0, see Theorem 3.1 in C-R). Our computations show that $\beta = \frac{1}{1.05}$ in our baseline parameterization suffices.

Result 3 *If organizational forgetting is absent ($\delta = 0$), then IID holds. Thus, ID holds.*

Even if an equilibrium satisfies ID and IID it is not clear that the industry is inevitably progressing towards monopolization. If the price gap between the firms is small, then the impact of ID and IID on industry structure and dynamics may be trivial.²⁹ In such a scenario, the leader charges a *slightly* lower price than the follower and this gap widens *a bit* over time. However, with even a modest degree of horizontal product differentiation, the firms still split sales more or less equally and thus move down the learning curve in tandem. This is exactly what happens when $\delta = 0$. For example, the flat equilibrium without well ($\rho = 0.85$, $\delta = 0$) satisfies IID and thus ID, but with a maximum expected Herfindahl index of 0.52 the industry is essentially a symmetric duopoly at all times. More generally, as Figure 3 shows, in the absence of organizational forgetting asymmetries are modest if they arise at all. Although ID and IID hold, the maximum expected Herfindahl index across all equilibria is 0.67 (attained at $\rho = 0.65$). Hence, ID and IID are not sufficient for economically meaningful market dominance.

ID and IID are also not necessary for market dominance. The extra-trenchy equilibrium ($\rho = 0.85$, $\delta = 0.08$), for example, violates ID and thus IID. Yet, the industry is likely to be a near-monopoly at all times. More generally, while the empirical studies of Argote et al. (1990), Darr et al. (1995), Benkard (2000), Shafer et al. (2001), and Thompson (2003) warrant accounting for organizational forgetting in a model of learning-by-doing, doing so may cause ID and IID to fail.

Proposition 8 *If organizational forgetting is present ($\delta > 0$), then IID may fail. Also ID may fail.*

In the absence of organizational forgetting, C-R have already shown that ID and IID may fail for intermediate values of β (see their Remark C.5 on p. 1136). Result 3 and Proposition 8 make the comparative dynamics point that ID and IID may hold when $\delta = 0$ but fail when $\delta > 0$ (holding fixed the remaining parameters). Figure 10 illustrates Proposition 8 by plotting the share of equilibria that violate IID (upper panel) and ID (lower panel). As can be seen, all equilibria fail to obey IID unless forgetting or learning is very weak. Even

²⁹Indeed, C-R show in their Theorem 3.2 that $p^*(\mathbf{e}) \rightarrow p^\dagger(m, m)$ for all $\mathbf{e} \in \{1, \dots, M\}^2$ as $\beta \rightarrow 1$, i.e., both firms price *as if* at the bottom of their learning curves. This suggests that the price gap may be small for “reasonable” discount factors.

	flat eqbm. without well	flat eqbm. with well	trenchy eqbm.	extra- trenchy eqbm.
leading example	$\rho = 0.85,$ $\delta = 0$	$\rho = 0.85,$ $\delta = 0.0275$	$\rho = 0.85,$ $\delta = 0.0275$	$\rho = 0.85,$ $\delta = 0.08$
preemption battle (well)	no	yes	no	no
price war triggered by imminent threat (diagonal trench)	no	no	yes	yes
price war triggered by distant threat (sideways trench)	no	no	no	yes
short-run market dominance	no	yes	yes	yes
long-run market dominance	no	no	yes, modest	yes, extreme
dominance properties	yes	no, mostly	no, mostly	no, mostly

Table 4: Pricing behavior and industry dynamics.

violations of ID are extremely common, especially for forgetting rates δ in the empirically relevant range below 0.1.

Of course, we do not argue that the concepts of ID and IID have no place in the analysis of industry dynamics. Caution, however, is advisable. Since ID and IID are neither necessary nor sufficient for market dominance, making inferences about the evolution of the industry on their basis alone may be misleading.

6.4 Summary

Table 4 summarizes the broad patterns of pricing behavior and industry dynamics. Acknowledging that the know-how gained through learning-by-doing can be lost through organizational forgetting is important. Generally speaking, organizational forgetting is associated with “trenchier” equilibria, more aggressive behavior, and more concentrated industries both in the short run and in the long run. Moreover, the dominance properties of firms’ pricing behavior can break down in the presence of organizational forgetting.

The key difference between a model with and without organizational forgetting is that in the former a firm can move both forward to a higher state and backward to a lower state. This possibility of bidirectional movement enhances the advantage-building and advantage-defending motives. By winning a sale, a firm makes itself less, and its rival more, vulnerable to organizational forgetting. This can create strong incentives to cut prices. Rather than impeding it, organizational forgetting therefore facilitates aggressive pricing as manifested in the trenchy and extra-trenchy equilibria that we have identified.

7 Organizational forgetting and multiple equilibria

While the equilibrium is unique if organizational forgetting is either absent ($\delta = 0$) or certain ($\delta = 1$), multiple equilibria are common for intermediate degrees of forgetting. Surprisingly, for some values of ρ and δ , the equilibria range from “peaceful coexistence” to “trench warfare.” Consequently, in addition to primitives of learning-by-doing and organizational forgetting, the equilibrium by itself is an important determinant of pricing behavior and industry dynamics.

Why do multiple equilibria arise in our model? To explore this question, think about the strategic situation faced by firms in setting prices in state \mathbf{e} . The value of continued play to firm n is given by the conditional expectation of its value function, $\bar{V}_{n1}(\mathbf{e})$ and $\bar{V}_{n2}(\mathbf{e})$, as defined in equations (2) and (3). Holding the value of continued play fixed, the strategic situation in state \mathbf{e} is akin to a static game. If the reaction functions in this game intersect more than once, then multiple equilibria arise. On the other hand, if they intersect only once irrespective of the value of continued play, then we say that a model satisfies *statewise uniqueness*.

Proposition 9 *Statewise uniqueness holds.*

Not surprisingly the proof of Proposition 9 relies on the functional form of demand. This is reminiscent of the restrictions on demand (e.g., log-concavity) that Caplin & Nalebuff (1991) set forth to guarantee uniqueness of Nash equilibrium in their analysis of static price-setting games.

Given that the model satisfies statewise uniqueness, multiple equilibria must arise from firms’ expectations regarding the value of continued play. To see this, consider again state \mathbf{e} . The intersection of the reaction functions constitutes a Nash equilibrium in prices in a subgame in which firm n believes that its value of continued play is given by $\bar{V}_{n1}(\mathbf{e})$ and $\bar{V}_{n2}(\mathbf{e})$. If firms have rational expectations, i.e., if the conjectured value of continued play is actually attained, then these prices constitute an equilibrium of our dynamic stochastic game. In our model, taking the value of continued play as given, the reaction functions intersect only once because we have statewise uniqueness, but there may be more than one value of continued play that is consistent with rational expectations. In this sense multiplicity is rooted in the dynamics of the model.

The key driver of multiplicity is organizational forgetting. Dynamic competition with learning and forgetting is like racing down an upward-moving escalator. Unless a firm makes sales at a rate that exceeds the rate at which it loses know-how through forgetting, its marginal cost is bound to increase. The inflow of know-how into the industry is one unit per period whereas in expectation the outflow in state \mathbf{e} is $\Delta(e_1) + \Delta(e_2)$. Consider

state (e, e) , where $e \geq m$, on the diagonal of the state space at or beyond the bottom of the learning curve. If $1 \ll 2\Delta(e)$, then it is impossible that both firms reach the bottom of their learning curves and remain there. Knowing this, firms have no choice but to price aggressively. The result is trench warfare as each firm uses price cuts to push the state to its side of the diagonal and keep it there. If, however, $1 \gg 2\Delta(e)$, then it is virtually inevitable that both firms reach the bottom of their learning curves, and firms may as well price softly. In both cases, the primitives of the model tie down the equilibrium.

This is no longer the case if $1 \approx 2\Delta(e)$, setting the stage for multiple equilibria as diverse as peaceful coexistence and trench warfare. If firms believe that they cannot peacefully coexist at the bottom of their learning curves and that one firm will come to dominate the market, then both firms will cut their prices in the hope of acquiring a competitive advantage early on and maintaining it throughout. This naturally leads to trench warfare and market dominance. If, however, firms believe that they can peacefully coexist at the bottom of their learning curves, then neither firm cuts its price. Soft pricing, in turn, ensures that the anticipated symmetric industry structure actually emerges. A back-of-the-envelope calculation is reassuring here. Recall that $m = 15$ and $M = 30$ in our baseline parameterization and observe that $1 = 2\Delta(15)$ implies $\delta \approx 0.045$, $1 = 2\Delta(20)$ implies $\delta \approx 0.034$, and $1 = 2\Delta(30)$ implies $\delta \approx 0.023$. This range of forgetting rates, for which the inflow of know-how approximately equals the outflow, is indeed where multiplicity prevails (see again Figure 2).

A sufficient condition for uniqueness of equilibrium in a dynamic stochastic game with a finite state space is that the model satisfies statewise uniqueness and the movements through the state space are unidirectional. Statewise uniqueness precludes players' actions from giving rise to multiple equilibria and unidirectional movements preclude their expectations from doing so. The proof of Proposition 3 illustrates the power of this sufficient condition. Specifically, if $\delta = 0$ in our game, then a firm can never move backward to a lower state. Hence, once the industry reaches state (M, M) , it remains there forever, so that the value of future play in state (M, M) coincides with the value of being in this state *ad infinitum*. In conjunction with statewise uniqueness, this uniquely determines the value of being in state (M, M) . Next consider states $(M - 1, M)$ and $(M, M - 1)$. The value of future play in states $(M - 1, M)$ and $(M, M - 1)$ depends on the value of being in state (M, M) . Statewise uniqueness ensures that firms' prices in states $(M - 1, M)$ and $(M, M - 1)$ as well as the value of being in these states are uniquely determined. Continuing to work backwards establishes that the equilibrium is unique.

8 Robustness checks

We have conducted extensive robustness checks regarding our specification of the discount factor, demand (product differentiation, outside good, and choke price), learning-by-doing, and entry and exit. In the interest of brevity, we confine ourselves here to pointing out if and how our results regarding aggressive pricing behavior (wells and trenches), market dominance (persistent asymmetries), and multiple equilibria change with the specification of the model. A more detailed discussion can be found in the Online Appendix along with some further checks (frequency of sales, organizational forgetting) that we omit here.

8.1 Discount factor

Two extreme cases merit discussion. First, as $\beta \rightarrow 0$ and firms become more myopic, the wells and trenches vanish and we obtain a flat equilibrium without well. In the limit of $\beta = 0$ equation (20) implies that the equilibrium of our dynamic stochastic game reduces to the static Nash equilibrium irrespective of the forgetting rate δ . Second, as $\beta \rightarrow 1$, the wells and trenches deepen: More patient firms have a stronger incentive to cut prices in the present in order to seize the leadership position in the future. In addition to the four typical cases in Figure 4 we obtain other types of equilibria with more complex patterns of trenches. The fact that high discount factors exacerbate the multiplicity problem is hardly surprising in light of the folk theorems for repeated games (Friedman 1971, Rubinstein 1979, Fudenberg & Maskin 1986).

8.2 Demand

Product differentiation. A higher degree of horizontal product differentiation σ lowers the extent to which firms interact strategically with each other. As $\sigma \rightarrow \infty$, firms eventually become monopolists and have no incentive to cut prices in order to acquire or defend a competitive advantage. As a result, the equilibrium is unique and the industry evolves symmetrically.

Outside good. As in C-R we assume that the buyer always purchases from one of the two firms in the industry. Allowing the buyer to instead choose an alternative made from a substitute technology (outside good) implies that the price elasticity of aggregate demand for the two competing firms is no longer zero. If the outside good is made sufficiently attractive, then in state (1, 1) the probability that either firm wins the one unit of demand that is available each period becomes small unless they price aggressively—below marginal cost—against the outside good. Making it down from the top of its learning curve consequently requires a firm to incur substantial losses in the near term. In the long term, however,

fighting one's way down the learning curve has substantial rewards because the outside good is a much less formidable competitor to a firm at the bottom of its learning curve than to a firm at the top.

If the discount factor is held fixed at its baseline value, then even a moderately attractive outside good sufficiently constrains firms' pricing behavior so that we no longer have sunspots for a progress ratio of $\rho = 1$. If we further increase the attractiveness of the outside good, then the rewards from fighting one's way down the learning curve become too far off in the future to justify the required aggressive pricing with its attendant near-term losses. In the ensuing equilibrium the price that a firm charges is fairly insensitive to its rival's stock of know-how because the outside good is the firm's main competitor and wins the sale most periods. As a result the inflow of know-how into the industry through learning is much smaller than the outflow through forgetting. This implies that the equilibrium is unique and entails both firms being stuck at the top of their learning curves. Trenchy and extra-trenchy equilibria disappear.

But if the discount factor is increased as the outside good is made increasingly attractive, then the near-term losses of fighting one's way down the learning curve do not overwhelm the long-term rewards from doing so. Firms price aggressively in trenchy and extra-trenchy equilibria. Therefore, provided the discount factor is sufficiently close to one, the presence of an economically significant price elasticity of aggregate demand does not seem to change the variety and multiplicity of equilibria in any fundamental way.

Choke price. As in C-R our logit specification for demand ensures that a firm always has a positive probability of making a sale and, in the absence of forgetting, must therefore eventually reach the bottom of its learning curve. This precludes long-run market dominance to occur in the absence of organizational forgetting.

Suppose instead that the probability that firm n makes a sale is given by a linear specification. Due to the choke price in the linear specification, a firm is able to surely deny its rival a sale by pricing sufficiently aggressively. Given a sufficiently low degree of horizontal product differentiation, firms at the top of their learning curves fight a preemption battle. The industry remains in an asymmetric structure as the winning firm takes advantage of the choke price to stall the losing firm at the top of its learning curve. In other words, the choke price is a shut-out model element that can lead to persistent asymmetries even in the absence of organizational forgetting.

8.3 Learning-by-doing

Following C-R we assume that $m < M$ represents the stock of know-how at which a firm reaches the bottom of its learning curve. In a *bottomless learning* specification with $m =$

M , we obtain other types of equilibria in addition to the four typical cases in Figure 4. Particularly striking is the *plateau equilibrium*. This equilibrium is similar to a trenchy equilibrium except that the diagonal trench is interrupted by a region of very soft price competition. On this plateau both firms charge prices well above cost. This “cooperative” behavior contrasts markedly with the price war of the diagonal trench.

8.4 Entry and exit

We assume that at any point in time there is a total of N firms, each of which can be either an incumbent firm or a potential entrant. Once an incumbent firm exits the industry, it perishes and a potential entrant automatically takes its “slot” and has to decide whether or not to enter.

Organizational forgetting remains a source of aggressive pricing behavior, market dominance, and multiple equilibria in the general model with entry and exit. The possibility of exit adds another component to the prize from winning a sale because, by winning a sale, a firm may move the industry to a state in which its rival is likely to exit. But if the rival exits, then it may be replaced by an entrant that comes into the industry at the top of its learning curve or it may not be replaced at all. As a result, pricing behavior is more aggressive than in the basic model without entry and exit. This leads to more pronounced asymmetries both in the short run and in the long run.

Because entry and exit are shut-out model elements, asymmetries can arise and persist even in the absence of organizational forgetting (see the Online Appendix for a concrete example). Entry and exit may also give rise to multiple equilibria as C-R have already shown (see their Theorem 4.1).

9 Conclusions

Learning-by-doing and organizational forgetting have been shown to be important in a variety of industrial settings. This paper provides a general model of dynamic competition that accounts for these economic fundamentals and shows how they shape industry structure and dynamics. We contribute to the numerical analysis of industry dynamics in two ways. First, we show that there are equilibria that the P-M algorithm cannot compute. Second, we propose a homotopy algorithm that allows us to describe in detail the structure of the set of solutions to our model.

In contrast to the present paper, the theoretical literature on learning-by-doing has largely ignored organizational forgetting. Moreover, it has mainly focused on the dominance properties of firms’ pricing behavior. By directly examining industry dynamics, we are able to show that ID and IID may not be sufficient for economically meaningful market

dominance. By generalizing the existing models of learning, we are able to show that these dominance properties break down with even a small degree of forgetting. Yet, it is precisely in the presence of organizational forgetting that market dominance ensues both in the short run and in the long run.

Our analysis of the role of organizational forgetting reveals that learning and forgetting are distinct economic forces. Forgetting, in particular, does not simply negate learning. The unique role played by organizational forgetting comes about because it makes bidirectional movements through the state space possible. As a consequence, a model with forgetting can give rise to aggressive pricing behavior, market dominance, and multiple equilibria, whereas a model without forgetting cannot.

Diagonal and sideways trenches are part and parcel to the self-reinforcing mechanisms that lead to market dominance. Since the leadership position is aggressively defended, firms fight a price war to attain it. This provides all the more reason to aggressively defend the leadership position because if it is lost, then another price war ensues. This seems like a good story to tell. Our computations show that this is not just an intuitively sensible story but also a logically consistent one that—perhaps—plays out in real markets.

Appendix

Proof of Proposition 1. *Part (i):* The basic differential equations (13) set

$$\delta'(s) = \det \left(\frac{\partial \mathbf{F}(\mathbf{x}(s); \delta(s), \rho)}{\partial \mathbf{x}} \right).$$

The Jacobian $\frac{\partial \mathbf{F}(\mathbf{x}(s); \delta(s), \rho)}{\partial \mathbf{x}}$ is a $(2M^2 \times 2M^2)$ matrix and therefore has an even number of eigenvalues. Its determinant is the product of its eigenvalues. Hence, if $\delta'(s) \leq 0$, then there exists at least one real nonnegative eigenvalue. (Suppose to the contrary that all eigenvalues are either complex or real and negative. Since the number of complex eigenvalues is even, so is the number of real eigenvalues. Moreover, the product of a conjugate pair of complex eigenvalues is positive, as is the product of an even number of real negative eigenvalues.)

To relate the P-M algorithm to our homotopy algorithm, let $(\mathbf{x}(s), \delta(s)) \in \mathbf{F}^{-1}(\rho)$ be a parametric path of equilibria. We show in the Online Appendix that

$$\left. \frac{\partial \mathbf{G}(\mathbf{x}(s))}{\partial \mathbf{x}} \right|_{(\delta(s), \rho)} = \frac{\partial \mathbf{F}(\mathbf{x}(s); \delta(s), \rho)}{\partial \mathbf{x}} + \mathbf{I}, \quad (21)$$

where \mathbf{I} denotes the $(2M^2 \times 2M^2)$ identity matrix.

The proof is completed by recalling a basic result from linear algebra: Let \mathbf{A} be an arbitrary matrix and $\zeta(\mathbf{A})$ its spectrum. Then $\zeta(\mathbf{A} + \mathbf{I}) = \zeta(\mathbf{A}) + 1$ (see Proposition A.17 in Appendix A of Bertsekas & Tsitsiklis 1997). Hence, because $\frac{\partial \mathbf{F}(\mathbf{x}(s); \delta(s), \rho)}{\partial \mathbf{x}}$ has at least one real nonnegative eigenvalue, it follows from equation (21) that $\left. \frac{\partial \mathbf{G}(\mathbf{x}(s))}{\partial \mathbf{x}} \right|_{(\delta(s), \rho)}$ has at least

one real eigenvalue equal to or bigger than unity. Hence, $\varrho \left(\left. \frac{\partial \mathbf{G}(\mathbf{x}(s))}{\partial \mathbf{x}} \right|_{(\delta(s), \rho)} \right) \geq 1$.

Part (ii): Consider the iteration $\mathbf{x}^{k+1} = \tilde{\mathbf{G}}(\mathbf{x}^k) = \omega \mathbf{G}(\mathbf{x}^k) + (1 - \omega) \mathbf{x}^k$, where $\omega > 0$. Using equation (21) its Jacobian at $(\mathbf{x}(s), \delta(s)) \in \mathbf{F}^{-1}(\rho)$ is

$$\left. \frac{\partial \tilde{\mathbf{G}}(\mathbf{x}(s))}{\partial \mathbf{x}} \right|_{(\delta(s), \rho)} = \omega \left. \frac{\partial \mathbf{G}(\mathbf{x}(s))}{\partial \mathbf{x}} \right|_{(\delta(s), \rho)} + (1 - \omega) \mathbf{I} = \omega \frac{\partial \mathbf{F}(\mathbf{x}(s); \delta(s), \rho)}{\partial \mathbf{x}} + \mathbf{I}.$$

As before it follows that $\varrho \left(\left. \frac{\partial \tilde{\mathbf{G}}(\mathbf{x}(s))}{\partial \mathbf{x}} \right|_{(\delta(s), \rho)} \right) \geq 1$. ■

Proof of Proposition 3. We rewrite the Bellman equations and FOCs in state \mathbf{e} as

$$V_1 = D_1(p_1, p_2) (p_1 - c(e_1) + \beta (\bar{V}_{11} - \bar{V}_{12})) + \beta \bar{V}_{12}, \quad (22)$$

$$V_2 = D_2(p_1, p_2) (p_2 - c(e_2) + \beta (\bar{V}_{22} - \bar{V}_{21})) + \beta \bar{V}_{21}, \quad (23)$$

$$0 = \frac{\sigma}{D_2(p_1, p_2)} - (p_1 - c(e_1) + \beta (\bar{V}_{11} - \bar{V}_{12})), \quad (24)$$

$$0 = \frac{\sigma}{D_1(p_1, p_2)} - (p_2 - c(e_2) + \beta (\bar{V}_{22} - \bar{V}_{21})), \quad (25)$$

where, to simplify the notation, V_n is shorthand for $V_n(\mathbf{e})$, \bar{V}_{nk} for $\bar{V}_{nk}(\mathbf{e})$, p_n for $p_n(\mathbf{e})$, etc. and we use the fact that $D_1(p_1, p_2) + D_2(p_1, p_2) = 1$.

Case (i): First suppose $\delta = 0$. The proof proceeds in a number of steps. In step 1, we establish that the equilibrium in state (M, M) is unique. In step 2a, we assume that there is a unique equilibrium in state $(e_1 + 1, M)$, where $e_1 \in \{1, \dots, M - 1\}$, and show that this implies that the equilibrium in state (e_1, M) is unique. In step 2b, we assume that there is a unique equilibrium in state $(M, e_2 + 1)$, where $e_2 \in \{1, \dots, M - 1\}$, and show that this implies that the equilibrium in state (M, e_2) is unique. By induction, steps 1, 2a, and 2b establish uniqueness along the upper edge of the state space. In step 3, we assume that there is a unique equilibrium in states $(e_1 + 1, e_2)$ and $(e_1, e_2 + 1)$, where $e_1 \in \{1, \dots, M - 1\}$ and $e_2 \in \{1, \dots, M - 1\}$, and show that this implies that the equilibrium in state (e_1, e_2) is unique. Hence, uniqueness in state $(M - 1, M - 1)$ follows from uniqueness in states $(M, M - 1)$ and $(M - 1, M)$, uniqueness in state $(M - 2, M - 1)$ from uniqueness in states $(M - 1, M - 1)$ and $(M - 2, M)$, etc. Working backwards gives uniqueness in states $(e_1, M - 1)$, where $e_1 \in \{1, \dots, M - 1\}$. This, in turn, gives uniqueness in states $(e_1, M - 2)$, where $e_1 \in \{1, \dots, M - 1\}$, etc.

Step 1: Consider state $\mathbf{e} = (M, M)$. From the definition of the state-to-state transitions in Section 2, we have

$$\bar{V}_{11} = \bar{V}_{12} = V_1, \quad \bar{V}_{21} = \bar{V}_{22} = V_2.$$

Imposing these restrictions and solving equations (22) and (23) for V_1 and V_2 , respectively,

yields

$$V_1 = \frac{D_1(p_1, p_2)(p_1 - c(e_1))}{1 - \beta}, \quad (26)$$

$$V_2 = \frac{D_2(p_1, p_2)(p_2 - c(e_2))}{1 - \beta}. \quad (27)$$

Simplifying equations (24) and (25) yields

$$0 = \frac{\sigma}{D_2(p_1, p_2)} - (p_1 - c(e_1)) = F_1(p_1, p_2), \quad (28)$$

$$0 = \frac{\sigma}{D_1(p_1, p_2)} - (p_2 - c(e_2)) = F_2(p_1, p_2). \quad (29)$$

The system of equations (28) and (29) determines equilibrium prices. Once we have established that there is a unique solution for p_1 and p_2 , equations (26) and (27) immediately ascertain that V_1 and V_2 are unique.

Let $p_1^{\sharp}(p_2)$ and $p_2^{\sharp}(p_1)$ be defined by

$$F_1(p_1^{\sharp}(p_2), p_2) = 0, \quad F_2(p_1, p_2^{\sharp}(p_1)) = 0$$

and set $F(p_1) = p_1 - p_1^{\sharp}(p_2^{\sharp}(p_1))$. The p_1 that solves the system of equations (28) and (29) is the solution to $F(p_1) = 0$, and this solution is unique provided that $F(p_1)$ is strictly monotone. The implicit function theorem yields

$$F'(p_1) = 1 - \frac{\left(-\frac{\partial F_1}{\partial p_2}\right)}{\frac{\partial F_1}{\partial p_1}} \frac{\left(-\frac{\partial F_2}{\partial p_1}\right)}{\frac{\partial F_2}{\partial p_2}}.$$

Straightforward differentiation shows that

$$\begin{aligned} \frac{\left(-\frac{\partial F_1}{\partial p_2}\right)}{\frac{\partial F_1}{\partial p_1}} &= \frac{-\frac{D_1(p_1, p_2)}{D_2(p_1, p_2)}}{-\frac{1}{D_2(p_1, p_2)}} = D_1(p_1, p_2) \in (0, 1), \\ \frac{\left(-\frac{\partial F_2}{\partial p_1}\right)}{\frac{\partial F_2}{\partial p_2}} &= \frac{-\frac{D_2(p_1, p_2)}{D_1(p_1, p_2)}}{-\frac{1}{D_1(p_1, p_2)}} = D_2(p_1, p_2) \in (0, 1). \end{aligned}$$

It follows that $F'(p_1) > 0$.

Step 2a: Consider state $\mathbf{e} = (e_1, M)$, where $e_1 \in \{1, \dots, M-1\}$. We have

$$\bar{V}_{12} = V_1, \quad \bar{V}_{22} = V_2.$$

Imposing these restrictions and solving equations (22) and (23) for V_1 and V_2 , respectively,

yields

$$V_1 = \frac{D_1(p_1, p_2)(p_1 - c(e_1) + \beta\bar{V}_{11})}{1 - \beta D_2(p_1, p_2)}, \quad (30)$$

$$V_2 = \frac{D_2(p_1, p_2)(p_2 - c(e_2) - \beta\bar{V}_{21}) + \beta\bar{V}_{21}}{1 - \beta D_2(p_1, p_2)}. \quad (31)$$

Substituting equations (30) and (31) into equations (24) and (25) and dividing through by $\frac{1-\beta}{1-\beta D_2(p_1, p_2)}$ and $\frac{1}{1-\beta D_2(p_1, p_2)}$, respectively, yields

$$0 = \frac{(1 - \beta D_2(p_1, p_2))\sigma}{(1 - \beta)D_2(p_1, p_2)} - (p_1 - c(e_1) + \beta\bar{V}_{11}) = G_1(p_1, p_2), \quad (32)$$

$$0 = \frac{(1 - \beta D_2(p_1, p_2))\sigma}{D_1(p_1, p_2)} - (p_2 - c(e_2) - \beta(1 - \beta)\bar{V}_{21}) = G_2(p_1, p_2). \quad (33)$$

The system of equations (32) and (33) determines equilibrium prices as a function of \bar{V}_{11} and \bar{V}_{21} . These are given by $V_1(e_1 + 1, M)$ and $V_2(e_1 + 1, M)$, respectively, and are unique by hypothesis. As in step 1, once we have established that there is a unique solution for p_1 and p_2 , equations (30) and (31) immediately ascertain that, in state $\mathbf{e} = (e_1, M)$, V_1 and V_2 are unique.

Proceeding as in step 1, set $G(p_1) = p_1 - p_1^{\natural}(p_2^{\natural}(p_1))$, where $p_1^{\natural}(p_2)$ and $p_2^{\natural}(p_1)$ are defined by $G_1(p_1^{\natural}(p_2), p_2) = 0$ and $G_2(p_1, p_2^{\natural}(p_1)) = 0$, respectively. We have to show that $G(\cdot)$ is strictly monotone. Straightforward differentiation shows that

$$\begin{aligned} \frac{\left(-\frac{\partial G_1}{\partial p_2}\right)}{\frac{\partial G_1}{\partial p_1}} &= \frac{-\frac{D_1(p_1, p_2)}{(1-\beta)D_2(p_1, p_2)}}{-\frac{1-\beta D_2(p_1, p_2)}{(1-\beta)D_2(p_1, p_2)}} = \frac{D_1(p_1, p_2)}{1 - \beta D_2(p_1, p_2)} \in (0, 1), \\ \frac{\left(-\frac{\partial G_2}{\partial p_1}\right)}{\frac{\partial G_2}{\partial p_2}} &= \frac{-\frac{(1-\beta)D_2(p_1, p_2)}{D_1(p_1, p_2)}}{-\frac{1-\beta D_2(p_1, p_2)}{D_1(p_1, p_2)}} = \frac{(1 - \beta)D_2(p_1, p_2)}{1 - \beta D_2(p_1, p_2)} \in (0, 1). \end{aligned}$$

It follows that $G'(p_1) > 0$.

Step 2b: Consider state $\mathbf{e} = (M, e_2)$, where $e_2 \in \{1, \dots, M - 1\}$. We have

$$\bar{V}_{11} = V_1, \quad \bar{V}_{21} = V_2.$$

The argument is completely symmetric to the argument in step 2a and therefore omitted.

Step 3: Consider state $\mathbf{e} = (e_1, e_2)$, where $e_1 \in \{1, \dots, M - 1\}$ and $e_2 \in \{1, \dots, M - 1\}$. The system of equations (24) and (25) determines equilibrium prices as a function of \bar{V}_{11} , \bar{V}_{12} , \bar{V}_{21} , and \bar{V}_{22} . These are given by $V_1(e_1 + 1, e_2)$, $V_1(e_1, e_2 + 1)$, $V_2(e_1 + 1, e_2)$, and $V_2(e_1, e_2 + 1)$, respectively, and are unique by hypothesis. As in step 1, once we have established that there is a unique solution for p_1 and p_2 , equations (22) and (23) immediately ascertain that, in state $\mathbf{e} = (e_1, e_2)$, V_1 and V_2 are unique.

Let $H_1(p_1, p_2)$ and $H_2(p_1, p_2)$ denote the RHS of equation (24) and (25), respectively. Proceeding as in step 1, set $H(p_1) = p_1 - p_1^{\natural}(p_2^{\natural}(p_1))$, where $p_1^{\natural}(p_2)$ and $p_2^{\natural}(p_1)$ are defined

by $H_1(p_1^\dagger(p_2), p_2) = 0$ and $H_2(p_1, p_2^\dagger(p_1)) = 0$, respectively. We have to show that $H(\cdot)$ is strictly monotone. Straightforward differentiation shows that

$$\begin{aligned} \left(\frac{-\frac{\partial H_1}{\partial p_2}}{\frac{\partial H_1}{\partial p_1}} \right) &= \frac{-\frac{D_1(p_1, p_2)}{D_2(p_1, p_2)}}{-\frac{1}{D_2(p_1, p_2)}} = D_1(p_1, p_2) \in (0, 1), \\ \left(\frac{-\frac{\partial H_2}{\partial p_1}}{\frac{\partial H_2}{\partial p_2}} \right) &= \frac{-\frac{D_2(p_1, p_2)}{D_1(p_1, p_2)}}{-\frac{1}{D_1(p_1, p_2)}} = D_2(p_1, p_2) \in (0, 1). \end{aligned}$$

It follows that $H'(p_1) > 0$.

Case (ii): Next suppose $\delta = 1$. A similar induction argument as in the case of $\delta = 0$ can be used to establish the claim except that in the case of $\delta = 1$ we anchor the argument in state $(1, 1)$ rather than state (M, M) . ■

Proof of Proposition 5. *Part (i):* Consider the static Nash equilibrium. The FOCs in state \mathbf{e} are

$$p_1^\dagger(\mathbf{e}) = c(e_1) + \frac{\sigma}{1 - D_1(p_1^\dagger(\mathbf{e}), p_2^\dagger(\mathbf{e}))}, \quad (34)$$

$$p_2^\dagger(\mathbf{e}) = c(e_2) + \frac{\sigma}{1 - D_2(p_1^\dagger(\mathbf{e}), p_2^\dagger(\mathbf{e}))}. \quad (35)$$

Equations (34) and (35) imply $p_1^\dagger(\mathbf{e}) > c(e_1)$ and $p_2^\dagger(\mathbf{e}) > c(e_2)$ and thus in particular $p^\dagger(m, m) > c(m)$. In addition, $p^\dagger(\mathbf{e}) = p^\dagger(m, m)$ because $c(e_1) = c(e_2) = c(m)$ for all $\mathbf{e} \in \{m, \dots, M\}^2$.

Turning to our dynamic stochastic game, suppose that $\delta = 0$. The proof of part (i) proceeds in a number of steps, similar to the proof of Proposition 3. In step 1, we establish that equilibrium prices in state (M, M) coincide with the static Nash equilibrium. In step 2a, we assume that the equilibrium in state $(e_1 + 1, M)$, where $e_1 \in \{m, \dots, M - 1\}$, coincides with the equilibrium in state (M, M) and show that this implies that the equilibrium in state (e_1, M) does the same. In step 2b, we assume that the equilibrium in state $(M, e_2 + 1)$, where $e_2 \in \{m, \dots, M - 1\}$, coincides with the equilibrium in state (M, M) and show that this implies that the equilibrium in state (M, e_2) does the same. In step 3, we assume that the equilibrium in states $(e_1 + 1, e_2)$ and $(e_1, e_2 + 1)$, where $e_1 \in \{m, \dots, M - 1\}$ and $e_2 \in \{m, \dots, M - 1\}$, coincides with the equilibrium in state (M, M) and show that this implies that the equilibrium in state (e_1, e_2) does the same. Also similar to the proof of Proposition 3, we continue to use V_n as shorthand for $V_n(\mathbf{e})$, \bar{V}_{nk} for $\bar{V}_{nk}(\mathbf{e})$, p_n for $p_n(\mathbf{e})$, etc.

Step 1: Consider state $\mathbf{e} = (M, M)$. From the proof of Proposition 3, equilibrium prices are determined by the system of equations (28) and (29). Since equations (28) and (29) are equivalent to equations (34) and (35), equilibrium prices are $p_1 = p_1^\dagger$ and $p_2 = p_2^\dagger$.

Substituting equation (28) into (26) and equation (29) into (27) yields equilibrium values

$$V_1 = \frac{\sigma D_1(p_1, p_2)}{(1 - \beta) D_2(p_1, p_2)}, \quad (36)$$

$$V_2 = \frac{\sigma D_2(p_1, p_2)}{(1 - \beta) D_1(p_1, p_2)}. \quad (37)$$

Step 2a: Consider state $\mathbf{e} = (e_1, M)$, where $e_1 \in \{m, \dots, M - 1\}$. Equilibrium prices are determined by the system of equations (32) and (33). Given $\bar{V}_{11} = V_1(e_1 + 1, M) = V_1(M, M)$ and $\bar{V}_{21} = V_2(e_1 + 1, M) = V_2(M, M)$, it is easy to see that, in state $\mathbf{e} = (e_1, M)$, $p_1 = p_1(M, M)$ and $p_2 = p_2(M, M)$ are a solution. Substituting equation (32) into (30) and equation (33) into (31) yields equilibrium values $V_1 = V_1(M, M)$ and $V_2 = V_2(M, M)$ as given by equations (36) and (37).

Step 2b: Consider state $\mathbf{e} = (M, e_2)$, where $e_2 \in \{m, \dots, M - 1\}$. The argument is completely symmetric to the argument in step 2a and therefore omitted.

Step 3: Consider state $\mathbf{e} = (e_1, e_2)$, where $e_1 \in \{m, \dots, M - 1\}$ and $e_2 \in \{m, \dots, M - 1\}$. Equilibrium prices are determined by the system of equations (24) and (25). Given $\bar{V}_{11} = V_1(e_1 + 1, e_2) = V_1(M, M)$, $\bar{V}_{12} = V_1(e_1, e_2 + 1) = V_1(M, M)$, $\bar{V}_{21} = V_2(e_1 + 1, e_2) = V_2(M, M)$, and $\bar{V}_{22} = V_2(e_1, e_2 + 1) = V_2(M, M)$, it is easy to see that, in state $\mathbf{e} = (e_1, e_2)$, $p_1 = p_1(M, M)$ and $p_2 = p_2(M, M)$ are a solution. Substituting equation (24) into (22) and equation (25) into (23) yields equilibrium values $V_1 = V_1(M, M)$ and $V_2 = V_2(M, M)$ as given by equations (36) and (37).

Part (ii): We show that $p_2(\mathbf{e}) > c(m)$ for all $e_1 \in \{1, \dots, m - 1\}$ and $e_2 \in \{m, \dots, M\}$. The claim follows because $p^*(\mathbf{e}) = p_2(\mathbf{e}^{[2]})$.

The proof of part (ii) proceeds in two steps. In step 1, we establish that the equilibrium price of firm 2 in state (e_1, M) , where $e_1 \in \{1, \dots, m - 1\}$, exceeds $c(m)$. In step 2, we extend the argument to states in which firm 2 has not yet reached the bottom of its learning curve. We proceed by induction: Assuming that the equilibrium in state $(e_1, e_2 + 1)$, where $e_1 \in \{1, \dots, m - 1\}$ and $e_2 \in \{m, \dots, M - 1\}$, coincides with the equilibrium in state (e_1, M) , we show that the equilibrium in state (e_1, e_2) does the same.

Step 1: Consider state $\mathbf{e} = (e_1, M)$, where $e_1 \in \{1, \dots, m - 1\}$. From the proof of Proposition 3, equilibrium prices are determined by the system of equations (32) and (33). In equilibrium we must have $V_n(\mathbf{e}) \geq 0$ for all $\mathbf{e} \in \{1, \dots, M\}^2$ because a firm can always set price equal to cost. Hence, $\bar{V}_{21} \geq 0$ and equation (33) implies $p_2 > c(m)$.

For reference in step 2 note that substituting equation (32) into (30) and equation (33) into (31) yields equilibrium values

$$V_1 = \frac{\sigma D_1(p_1, p_2)}{(1 - \beta) D_2(p_1, p_2)}, \quad (38)$$

$$V_2 = \frac{\sigma D_2(p_1, p_2) + \beta D_1(p_1, p_2) \bar{V}_{21}}{D_1(p_1, p_2)}. \quad (39)$$

Step 2: Consider state $\mathbf{e} = (e_1, e_2)$, where $e_1 \in \{1, \dots, m - 1\}$ and $e_2 \in \{m, \dots, M - 1\}$. Equilibrium prices are determined by the system of equations (24) and (25). Assuming $\bar{V}_{12} = V_1(e_1, e_2 + 1) = V_1(e_1, M)$ and $\bar{V}_{22} = V_2(e_1, e_2 + 1) = V_2(e_1, M)$ as given by equations

(38) and (39) in step 1, equations (24) and (25) collapse to equations (32) and (33). Hence, in state $\mathbf{e} = (e_1, e_2)$, $p_1 = p_1(e_1, M)$ and $p_2 = p_2(e_1, M)$ are a solution. Further substituting equation (24) into (22) and equation (25) into (23) yields equilibrium values $V_1 = V_1(e_1, M)$ and $V_2 = V_2(e_1, M)$ as given by equations (38) and (39). ■

Proof of Proposition 9. We rewrite the FOCs in state \mathbf{e} as

$$0 = \frac{\sigma}{D_2(p_1, p_2)} - (p_1 - c(e_1) + \beta (\bar{V}_{11} - \bar{V}_{12})), \quad (40)$$

$$0 = \frac{\sigma}{D_1(p_1, p_2)} - (p_2 - c(e_2) + \beta (\bar{V}_{22} - \bar{V}_{21})), \quad (41)$$

where, to simplify the notation, \bar{V}_{nk} is shorthand for $\bar{V}_{nk}(\mathbf{e})$, p_n for $p_n(\mathbf{e})$, etc. and we use the fact that $D_1(p_1, p_2) + D_2(p_1, p_2) = 1$. The system of equations (40) and (41) determines equilibrium prices. We have to establish that there is a unique solution for p_1 and p_2 irrespective of \bar{V}_{11} , \bar{V}_{12} , \bar{V}_{21} , and \bar{V}_{22} .

Let $H_1(p_1, p_2)$ and $H_2(p_1, p_2)$ denote the RHS of equation (40) and (41), respectively. Proceeding as in step 3 of the proof of Proposition 3, set $H(p_1) = p_1 - p_1^\sharp(p_2^\sharp(p_1))$, where $p_1^\sharp(p_2)$ and $p_2^\sharp(p_1)$ are defined by $H_1(p_1^\sharp(p_2), p_2) = 0$ and $H_2(p_1, p_2^\sharp(p_1)) = 0$, respectively. We have to show that $H(\cdot)$ is strictly monotone. Straightforward differentiation shows that

$$\begin{aligned} \left(-\frac{\partial H_1}{\partial p_2} \right) &= \frac{D_1(p_1, p_2)}{D_2(p_1, p_2)} = D_1(p_1, p_2) \in (0, 1), \\ \frac{\partial H_1}{\partial p_1} &= \frac{1}{D_2(p_1, p_2)} \\ \left(-\frac{\partial H_2}{\partial p_1} \right) &= \frac{D_2(p_1, p_2)}{D_1(p_1, p_2)} = D_2(p_1, p_2) \in (0, 1). \\ \frac{\partial H_2}{\partial p_2} &= \frac{1}{D_1(p_1, p_2)} \end{aligned}$$

It follows that $H'(p_1) > 0$. ■

References

- Akerberg, D., Benkard, L., Berry, S. & Pakes, A. (2007), Econometric tools for analyzing market outcomes, *in* J. Heckman & E. Leamer, eds, ‘Handbook of Econometrics’, Vol. 6A, North-Holland, Amsterdam, pp. 4171–4276.
- Alchian, A. (1963), ‘Reliability of progress curves in airframe production’, *Econometrica* **31**(4), 679–693.
- Argote, L., Beckman, S. & Epple, D. (1990), ‘The persistence and transfer of learning in an industrial setting’, *Management Science* **36**(2), 140–154.
- Athey, S. & Schmutzler, A. (2001), ‘Investment and market dominance’, *Rand Journal of Economics* **32**(1), 1–26.
- Bailey, C. (1989), ‘Forgetting and the learning curve: A laboratory study’, *Management Science* **35**(3), 340–352.

- Baloff, N. (1971), 'Extension of the learning curve: Some empirical results', *Operational Reserach Quarterly* **22**(4), 329–340.
- Benkard, L. (2000), 'Learning and forgetting: The dynamics of aircraft production', *American Economic Review* **90**(4), 1034–1054.
- Benkard, L. (2004), 'A dynamic analysis of the market for wide-bodied commercial aircraft', *Review of Economic Studies* **71**(3), 581–611.
- Berry, S. & Pakes, A. (2007), 'The pure characteristics demand model', *International Economic Review* **48**(4), 1193–1225.
- Bertsekas, D. & Tsitsiklis, J. (1997), *Parallel and distributed computation: Numerical methods*, Athena Scientific, Belmont.
- Bohn, R. (1995), 'Noise and learning in semiconductor manufacturing', *Management Science* **41**(1), 31–42.
- Borkovsky, R., Doraszelski, U. & Kryukov, S. (2008), A user's guide to solving dynamic stochastic games using the homotopy method, Working paper, Northwestern University, Evanston.
- Budd, C., Harris, C. & Vickers, J. (1993), 'A model of the evolution of duopoly: Does the asymmetry between firms tend to increase or decrease?', *Review of Economic Studies* **60**(3), 543–573.
- Cabral, L. & Riordan, M. (1994), 'The learning curve, market dominance, and predatory pricing', *Econometrica* **62**(5), 1115–1140.
- Cabral, L. & Riordan, M. (1997), 'The learning curve, predation, antitrust, and welfare', *Journal of Industrial Economics* **45**(2), 155–169.
- Caplin, A. & Nalebuff, B. (1991), 'Aggregation and imperfect competition: On the existence of equilibrium', *Econometrica* **59**(1), 26–59.
- Caves, R. & Porter, M. (1977), 'From entry barriers to mobility barriers: Conjectural decisions and contrived deterrence to new competition', *Quarterly Journal of Economics* **91**(2), 241–261.
- Darr, E., Argote, L. & Epple, D. (1995), 'The acquisition, transfer, and depreciation of knowledge in service organizations: Productivity in franchises', *Management Science* **41**(11), 1750–1762.
- Dasgupta, P. & Stiglitz, J. (1988), 'Learning-by-doing, market structure and industrial and trade policies', *Oxford Economics Papers* **40**(2), 246–268.
- DeJong, J. (1957), 'The effects of increasing skills on cycle time and its consequences for time standards', *Ergonomics* **1**(1), 51–60.
- Demichelis, S. & Germano, F. (2002), 'On (un)knots and dynamics in games', *Games and Economic Behavior* **41**, 46–60.

- Doraszelski, U. & Judd, K. (2004), Avoiding the curse of dimensionality in dynamic stochastic games, Technical working paper no. 304, NBER, Cambridge.
- Doraszelski, U. & Satterthwaite, M. (2008), Computable Markov-perfect industry dynamics, Working paper, Harvard University, Cambridge.
- Dudley, L. (1972), 'Learning and productivity changes in metal products', *American Economic Review* **62**(4), 662–669.
- Dutton, J. & Thomas, A. (1984), 'Treating progress functions as a managerial opportunity', *Academy of Management Review* **9**(2), 235–247.
- Ericson, R. & Pakes, A. (1995), 'Markov-perfect industry dynamics: A framework for empirical work', *Review of Economic Studies* **62**, 53–82.
- Friedman, J. (1971), 'A non-cooperative equilibrium for supergames', *Review of Economic Studies* **38**(1), 1–12.
- Fudenberg, D. & Maskin, E. (1986), 'The folk theorem in repeated games with discounting or with incomplete information', *Econometrica* **54**(3), 533–554.
- Fudenberg, D. & Tirole, J. (1983), 'Learning-by-doing and market performance', *Bell Journal of Economics* **14**(2), 522–530.
- Fudenberg, D. & Tirole, J. (1991), *Game theory*, MIT Press, Cambridge.
- Garcia, C. & Zangwill, W. (1979), 'An approach to homotopy and degree theory', *Mathematics of Operations Research* **4**(4), 390–405.
- Ghemawat, P. & Spence, M. (1985), 'Learning curve spillovers and market performance', *Quarterly Journal of Economics* **100**(Supplement), 839–852.
- Grieco, P. (2008), Numerical approximation of an equilibrium correspondence: Method and applications, Working paper, Northwestern University, Evanston.
- Gruber, H. (1992), 'The learning curve in the production of semiconductor memory chips', *Applied Economics* **24**, 885–895.
- Hatch, N. & Mowery, D. (1998), 'Process innovation and learning by doing in semiconductor manufacturing', *Management Science* **44**(1), 1461–1477.
- Herings, J. & Peeters, R. (2004), 'Stationary equilibria in stochastic games: Structure, selection, and computation', *Journal of Economic Theory* **118**, 32–60.
- Hirschmann, W. (1964), 'Profit from the learning curve', *Harvard Business Review* **42**(1), 125–139.
- Irwin, D. & Klenow, P. (1994), 'Learning-by-doing spillovers in the semiconductor industry', *Journal of Political Economy* **102**(6), 1200–1227.
- Jarmin, R. (1994), 'Learning by doing and competition in the early rayon industry', *Rand Journal of Economics* **25**(3), 441–454.

- Judd, K. (1998), *Numerical methods in economics*, MIT Press, Cambridge.
- Judd, K. & Schmedders, K. (2004), A computational approach to proving uniqueness in dynamic games, Working paper, Hoover Institution, Stanford.
- Kilbridge, M. (1962), 'A model of industrial learning', *Management Science* **8**(4), 516–527.
- Levy, F. (1965), 'Adaptation in the production process', *Management Science* **11**(6), 136–154.
- Lieberman, M. (1984), 'The learning curve and pricing in the chemical processing industries', *Rand Journal of Economics* **15**(2), 213–228.
- Maskin, E. & Tirole, J. (2001), 'Markov perfect equilibrium, I: Observable actions', *Journal of Economic Theory* **100**, 191–219.
- Pakes, A. (2008), Theorey and empirical work on imperfectly competitive markets, Fisher-Schultz lecture, Harvard University, Cambridge.
- Pakes, A. & McGuire, P. (1994), 'Computing Markov-perfect Nash equilibria: Numerical implications of a dynamic differentiated product model', *Rand Journal of Economics* **25**(4), 555–589.
- Pisano, G. (1994), 'Knowledge, integration, and the locus of learning: An empirical analysis of process development', *Strategic Management Journal* **15**, 85–100.
- Preston, L. & Keachie, E. (1964), 'Cost functions and progress functions: An integration', *American Economic Review* **54**(2), 100–107.
- Ross, D. (1986), 'Learning to dominate', *Journal of Industrial Economics* **34**(4), 337–353.
- Rubinstein, A. (1979), 'Equilibrium in supergames with the overtaking criterion', *Journal of Economic Theory* **21**(1), 1–9.
- Schmedders, K. (1998), 'Computing equilibria in the general equilibrium model with incomplete asset markets', *Journal of Economic Dynamics and Control* **22**, 1375–1401.
- Schmedders, K. (1999), 'A homotopy algorithm and an index theorem for the general equilibrium model with incomplete asset markets', *Journal of Mathematical Economics* **32**(2), 225–241.
- Shafer, S., Nembhard, D. & Uzumeri, M. (2001), 'The effects of worker learning, forgetting, and heterogeneity on assembly line productivity', *Management Science* **47**(12), 1639–1653.
- Shapley, L. (1953), 'Stochastic games', *Proceedings of the National Academy of Sciences* **39**, 1095–1100.
- Spence, M. (1981), 'The learning curve and competition', *Bell Journal of Economics* **12**(1), 49–70.

- Thompson, P. (2001), ‘How much did the Liberty shipbuilders learn? New evidence for an old case study’, *Journal of Political Economy* **109**(1), 103–137.
- Thompson, P. (2003), How much did the Liberty shipbuilders forget?, Working paper, Florida International University, Miami.
- Thornton, R. & Thompson, P. (2001), ‘Learning from experience and learning from others: An exploration of learning and spillovers in wartime shipbuilding’, *American Economic Review* **91**(5), 1350–1368.
- Vickers, J. (1986), ‘The evolution of market structure when there is a sequence of innovations’, *Journal of Industrial Economics* **35**(1), 1–12.
- Vives, X. (1999), *Oligopoly pricing: Old ideas and new tools*, MIT Press, Cambridge.
- Watson, L., Billups, S. & Morgan, A. (1987), ‘HOMPACK: A suite of codes for globally convergent homotopy algorithms’, *ACM Transactions on Mathematical Software* **13**(3), 281–310.
- Watson, L., Sosonkina, M., Melville, R., Morgan, A. & Walker, H. (1997), ‘Algorithm 777: HOMPACK90: A suite of Fortran 90 codes for globally convergent homotopy algorithms’, *ACM Transactions on Mathematical Software* **23**(4), 514–549.
- Wixted, M. & Ebbesen, E. (1991), ‘On the form of forgetting’, *Psychological Science* **2**(6), 409–415.
- Wright, T. (1936), ‘Factors affecting the cost of airplanes’, *Journal of Aeronautical Sciences* **3**(4), 122–128.
- Zangwill, W. & Garcia, C. (1981), *Pathways to solutions, fixed points, and equilibria*, Prentice Hall, Englewood Cliffs.
- Zimmerman, M. (1982), ‘Learning effects and the commercialization of new energy technologies: The case of nuclear power’, *Bell Journal of Economics* **13**(2), 297–310.

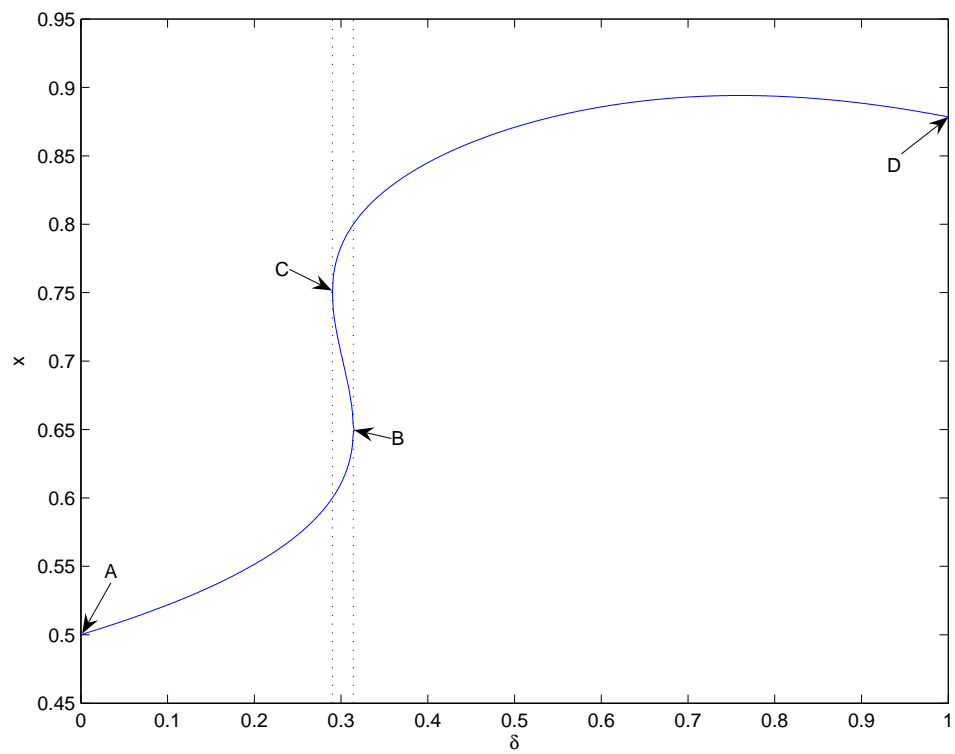


Figure 1: Homotopy example.

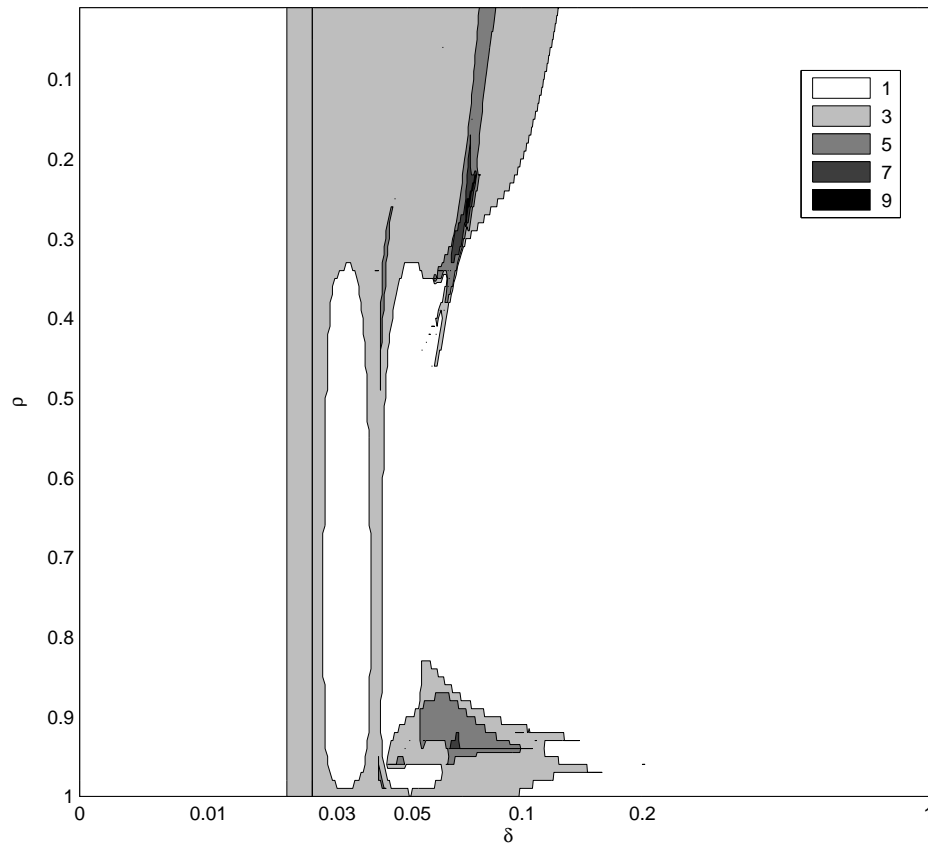


Figure 2: Number of equilibria.

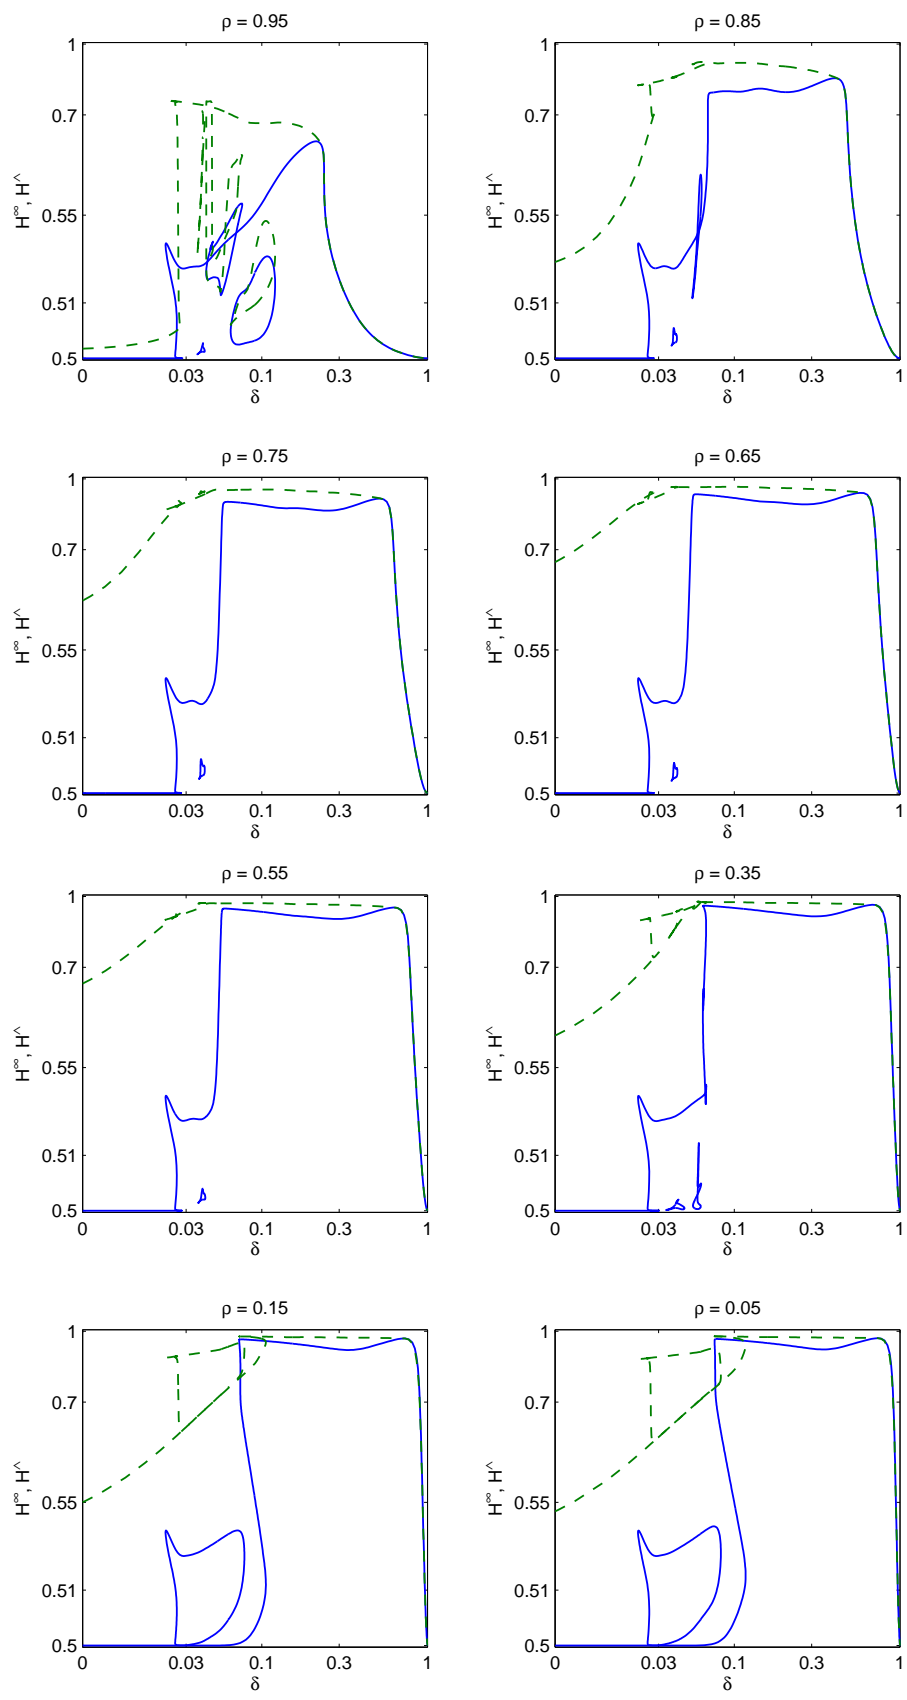


Figure 3: Limiting expected Herfindahl index H^∞ (solid line) and maximum expected Herfindahl index H^\wedge (dashed line).

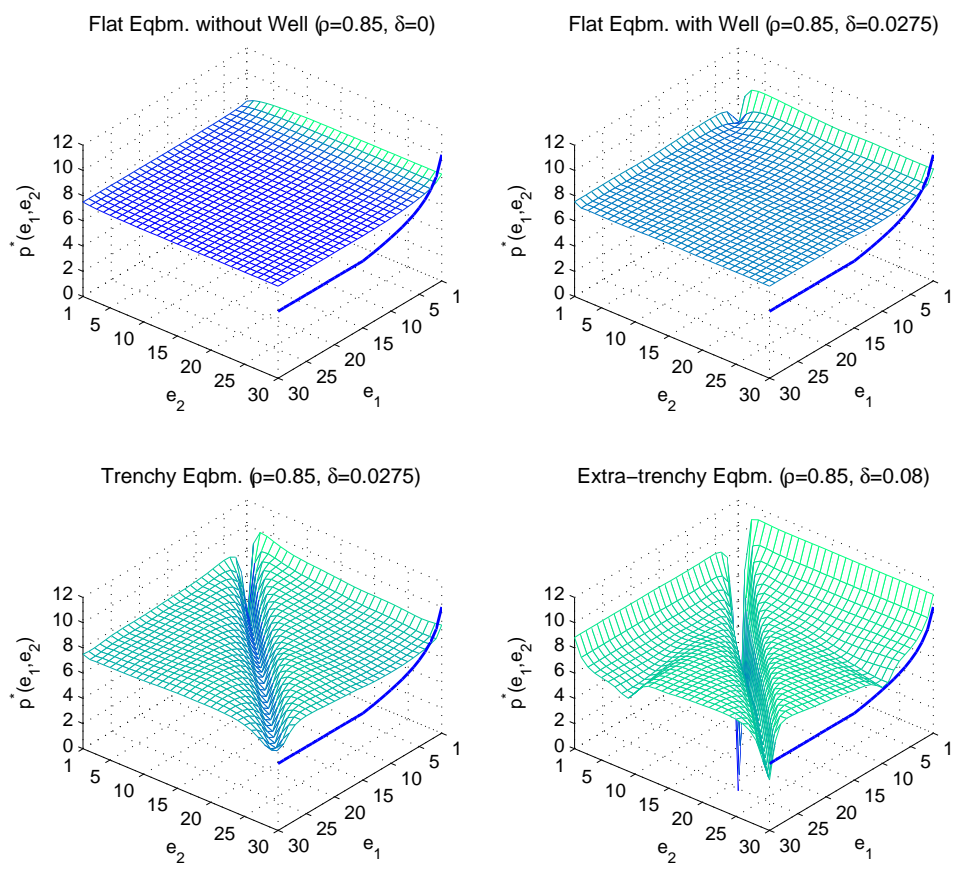


Figure 4: Policy function $p^*(e_1, e_2)$. Marginal cost $c(e_1)$ (solid line in $e_2 = 30$ -plane).

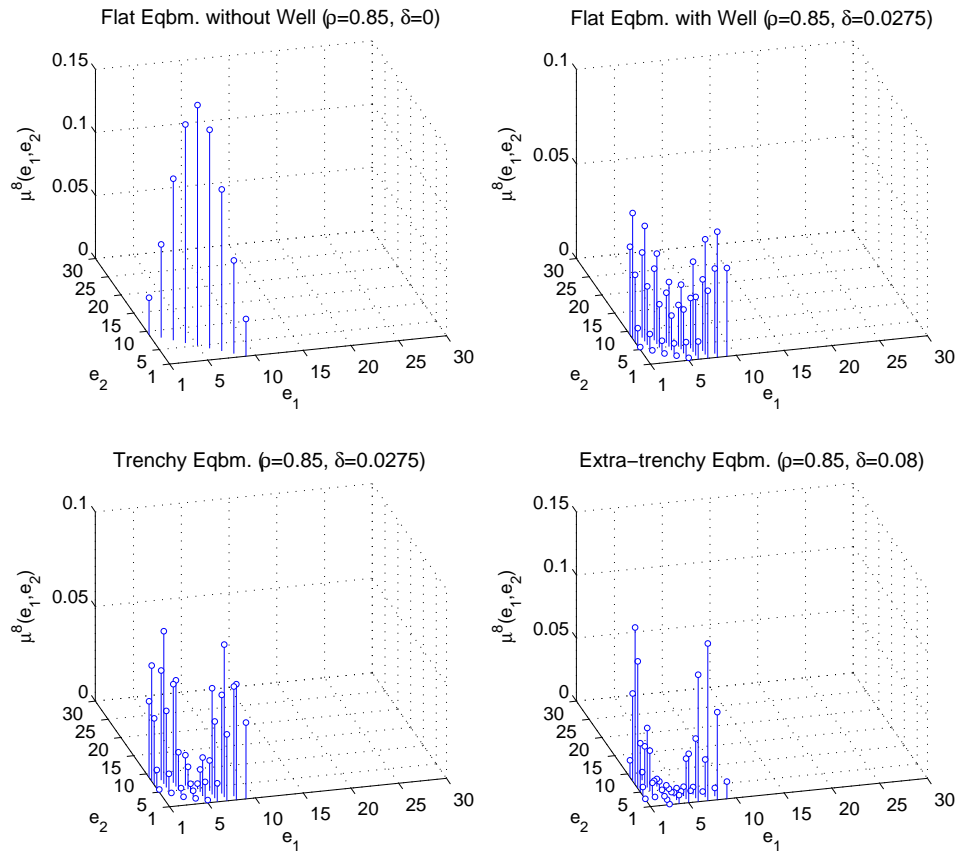


Figure 5: Transient distribution over states in period 8 given initial state $(1, 1)$.

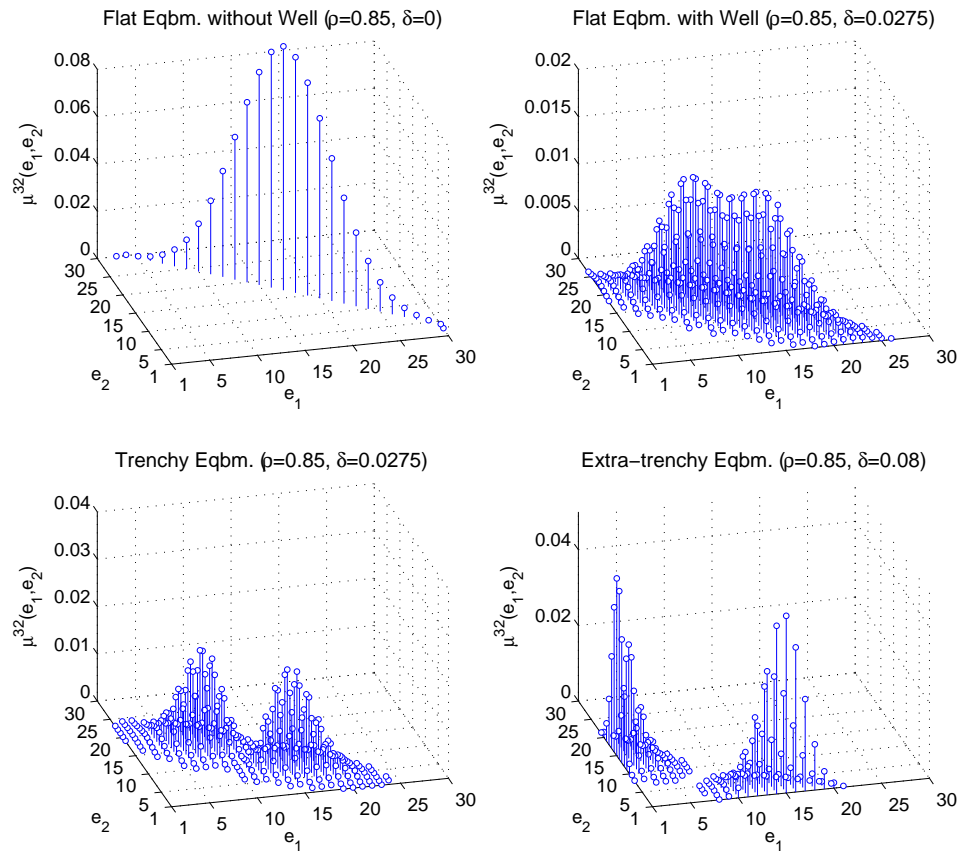


Figure 6: Transient distribution over states in period 32 given initial state $(1, 1)$.

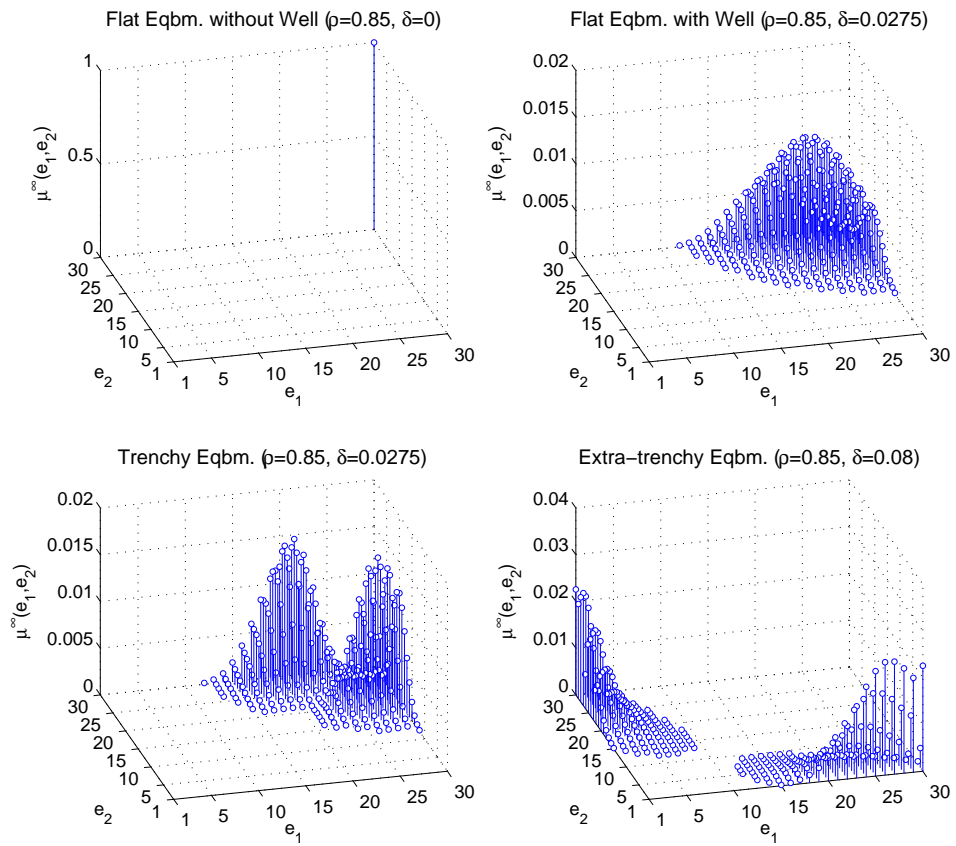


Figure 7: Limiting distribution over states.

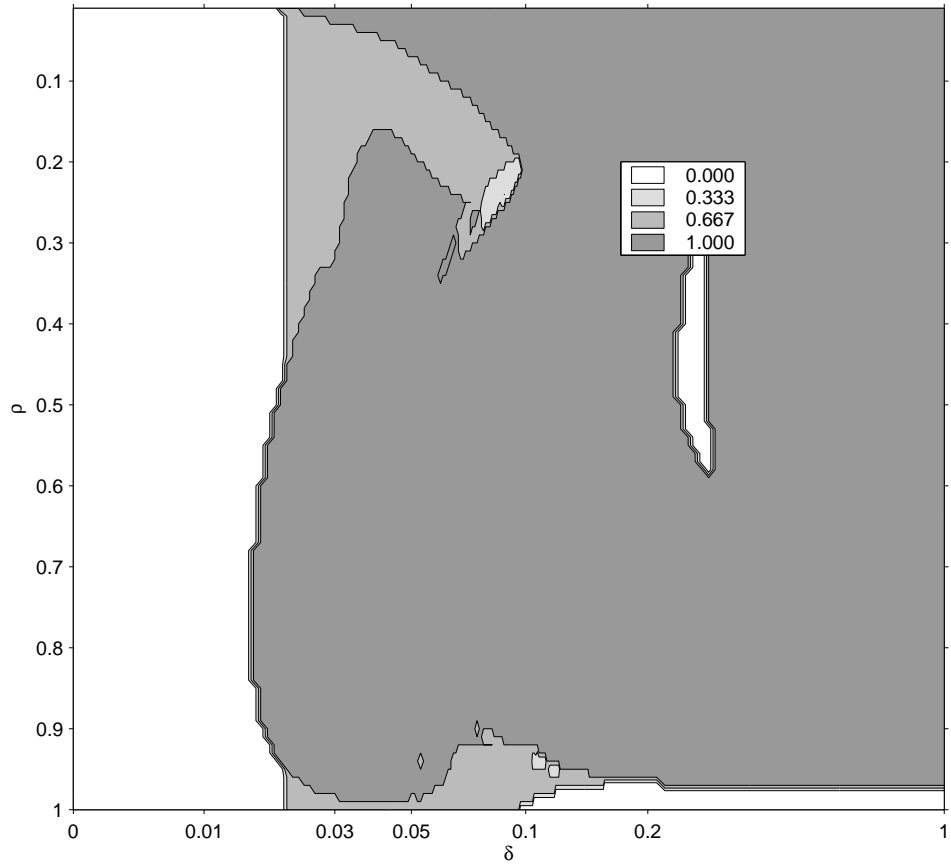


Figure 8: Share of equilibria with $p^*(\mathbf{e}) > p^\dagger(\mathbf{e})$ for some $\mathbf{e} \in \{1, \dots, M\}^2$.

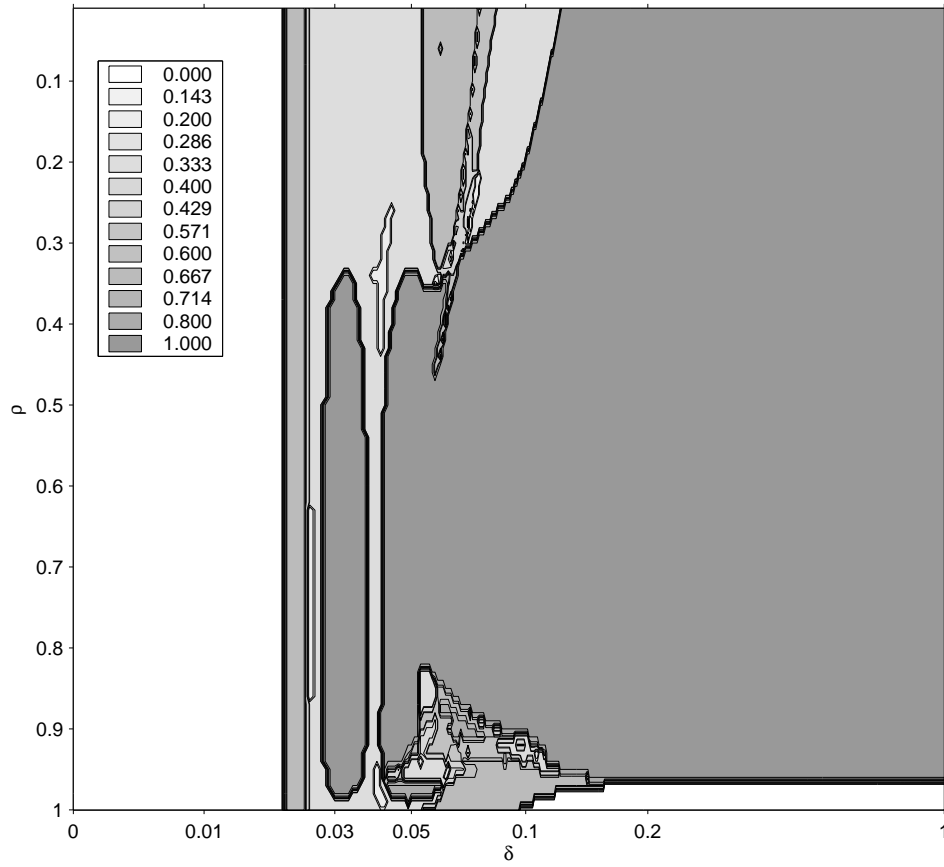


Figure 9: Share of equilibria with $p^*(\mathbf{e}) \leq c(m)$ for some $e_1 \in \{m, \dots, M\}$ and $e_2 \in \{1, \dots, M\}$.

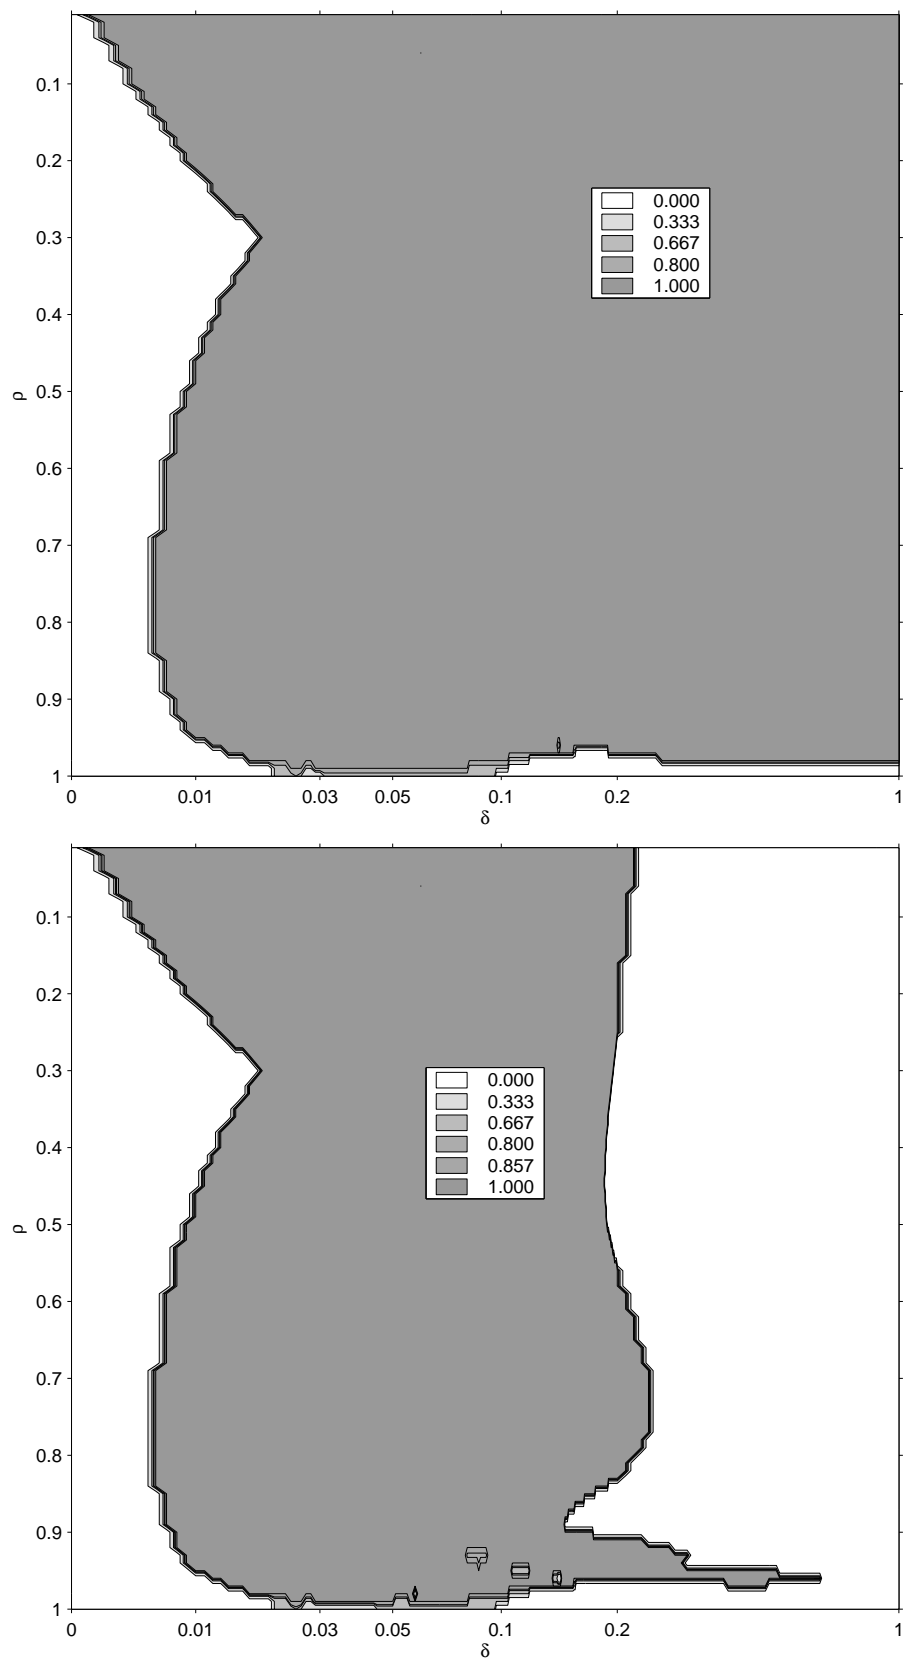


Figure 10: Share of equilibria violating IID (upper panel) and share of equilibria violating ID (lower panel).