

How Efficient Is Dynamic Competition? The Case of Price as Investment[†]

By DAVID BESANKO, ULRICH DORASZELSKI, AND YAROSLAV KRYUKOV*

We study industries where the price that a firm sets serves as an investment into lower cost or higher demand. We assess the welfare implications of the ensuing competition for the market using analytical and numerical approaches to compare the equilibria of a learning-by-doing model to the first-best planner solution. We show that dynamic competition leads to low deadweight loss. This cannot be attributed to similarity between the equilibria and the planner solution. Instead, we show how learning-by-doing causes the various contributions to deadweight loss to either be small or partly offset each other. (JEL D21, D25, D43, D83, L13)

In many “new economy” industries (e.g., Amazon and Barnes & Noble in e-book readers) and “old economy” industries (e.g., Boeing and Airbus in aircraft, Intel and AMD in microprocessors), the price that a firm sets plays the role of an investment. The investment role arises when the firm’s current price affects its future competitive position vis-à-vis its rivals. Examples include competition to accumulate production experience on a learning curve or to acquire a customer base in markets with network effects or switching costs. In these settings, a firm’s current sales translate into lower cost or higher demand in the future, and the firm can thus shape the evolution of the industry by pricing aggressively. Competition is dynamic as firms jostle for competitive advantage through the prices they set.

It is well understood that the investment role of price opens up a second dimension of competition between firms, namely competition for the market. At the same time, several high-profile antitrust cases, such as *United States v. Microsoft* in the late 1990s and *European Union v. Google* initiated in 2015 have taken aim at the market leaders in industries where the investment role of price is presumably important. This raises a question that is much less well understood: does unfettered

* Besanko: Kellogg School of Management, Northwestern University, Evanston, IL 60208 (email: d-besanko@kellogg.northwestern.edu); Doraszelski: Wharton School, University of Pennsylvania, Philadelphia, PA 19104 (email: doraszelski@wharton.upenn.edu); Kryukov: University of Pittsburgh Medical Center, Pittsburgh, PA 15219 (email: yaro.kryukov@gmail.com). Liran Einav was the coeditor for this article. We thank three anonymous referees, Guy Arie, Simon Board, Joe Harrington, Bruno Jullien, Ariel Pakes, Robert Porter, Michael Raith, Eric Rasmusen, Mike Riordan, Juuso Välimäki, and Mike Whinston for helpful discussions and suggestions. We also thank participants at seminars at Carnegie Mellon University, DG Comp at the European Commission, Federal Trade Commission Bureau of Economics, Helsinki Center for Economic Research, Johns Hopkins University, Toulouse School of Economics, Tulane University, and University of Rochester for their useful questions and comments. Alexandre Limonov provided outstanding research assistance. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

[†] Go to <https://doi.org/10.1257/aer.20180131> to visit the article page for additional materials and author disclosure statements.

competition for the market when price serves as an investment warrant regulatory scrutiny? The answer depends on how efficient dynamic competition is. If it is not very efficient and leads to high deadweight loss, then there are potential gains from intervening in the industry. However, if dynamic competition is very efficient, then the upside from regulatory scrutiny is limited, barring explicitly anti-competitive behavior.

It may seem intuitive that dynamic competition when price serves as an investment is fairly efficient. In contrast to rent-seeking models (Posner 1975) in which firms compete for market dominance by engaging in socially wasteful activities (e.g., lobbying), competition for the market when price serves as an investment occurs through low prices that transfer surplus to consumers. And as a by-product of these low prices, firms generate socially valuable learning economies in production or demand-side economies of scale.

However, this intuition is incomplete. First, prices that are too low relative to marginal cost can cause deadweight loss from overproduction. Second, dynamic competition may give rise to equilibria that entail predation-like behavior and monopolization of the industry in the long run (Dasgupta and Stiglitz 1988; Cabral and Riordan 1994; Athey and Schmutzler 2001; Besanko et al. 2010; Besanko, Doraszelski, and Kryukov 2014), causing welfare losses from monopoly pricing and suboptimal product variety. Third, as firms attempt to balance the gains from possibly monopolizing their industry against the losses from pricing aggressively as they jostle for competitive advantage, dynamic competition can distort entry behavior and result in coordination failures similar to the ones in natural monopoly markets highlighted by Bolton and Farrell (1990). It can also distort exit behavior if firms are reluctant to exit an industry destined to be monopolized and become caught up in wars of attrition (Smith 1974, Tirole 1988, Bulow and Klemperer 1999).

In sum, dynamic competition when price serves as an investment may cause deadweight losses through a number of distinct channels, and it is not a priori obvious what the magnitude of these deadweight losses may be. This paper makes a first attempt to assess how efficient dynamic competition is. We do so in a model of learning-by-doing along the lines of Cabral and Riordan (1994); Besanko et al. (2010); and Besanko, Doraszelski, and Kryukov (2014) that involves price competition in a differentiated products market with entry and exit. Learning-by-doing is important in a broad set of industries (see the references in Besanko et al. 2010).

We compare the Markov perfect equilibria of our learning-by-doing model to the solution of a first-best problem that has a social planner controlling pricing, entry, and exit decisions. Deadweight loss is the difference in the expected net present value of total surplus. We carefully explore the parameter space and multiple equilibria using the homotopy or path-following method in Besanko et al. (2010) and Besanko, Doraszelski, and Kryukov (2014). We show that dynamic competition does indeed tend to lead to low deadweight loss. It is less than 10 percent of the maximum value added by the industry in more than 65 percent of parameterizations and less than 20 percent of value added in more than 90 percent of parameterizations.¹ Moreover, the investment role of price tends to be socially beneficial; if

¹The maximum value added by the industry is the difference in the expected net present value of total surplus of the first-best planner solution and an industry that remains empty forever (see Section IIB).

we shut it down, then deadweight loss increases, often substantially, in more than 80 percent of parameterizations.

The low deadweight loss does not arise because equilibrium behavior and its implied industry dynamics are similar to the first-best planner solution. On the contrary, especially the “best” equilibria that entail the lowest deadweight loss are remarkably different from the first-best planner solution.

To identify the mechanisms through which dynamic competition leads to low deadweight loss, we decompose deadweight loss into a pricing distortion that captures differences in equilibrium pricing behavior from the first-best planner solution and a remaining non-pricing distortion that subsumes a suboptimal number of firms and products being offered by these firms, a suboptimal exploitation of learning economies, and cost-inefficient exit. The low deadweight loss boils down to three regularities in these components.

First, while the pricing distortion tends to be the largest contributor to deadweight loss, in the best equilibria it is quite low and it is “not so bad” in the “worst” equilibria that entail the highest deadweight loss. Analytical bounds on the pricing distortion reveal that this regularity is rooted in learning-by-doing itself. Second, the non-pricing distortion in the worst equilibria tends to be low because competition for the market resolves itself quickly and winnows out firms in a fairly efficient way. Third, in the best equilibria the non-pricing distortion tends to be higher. Dynamic competition tends to lead to over-entry and under-exit. While this has social costs (setup costs and forgone scrap values), it also has offsetting social benefits from additional product variety. Learning economies can be shown to accentuate these benefits by making the additional product variety less costly to procure.

All in all, we show that dynamic competition when price serves as an investment works remarkably well, not because competition for the market is a “magic bullet” that achieves full efficiency, but because the various contributions to deadweight loss either are small or partly offset each other. And this, in turn, happens because of learning-by-doing itself.

Our paper is related to a large literature on dynamic competition in settings where price plays the role of an investment. Besides the aforementioned models of learning-by-doing, this includes the models of network effects in Mitchell and Skrzypacz (2006); Chen, Doraszelski, and Harrington (2009); Dubé, Hitsch, and Chintagunta (2010); and Cabral (2011); switching costs in Dubé, Hitsch, and Rossi (2009) and Chen (2011); experience goods in Bergemann and Välimäki (1996) and Ching (2010); and habit formation in Bergemann and Välimäki (2006). As we show in Besanko, Doraszelski, and Kryukov (2014), these models share a key feature with our learning-by-doing model in that a firm has two distinct motives for pricing aggressively: to acquire competitive advantage or overcome competitive disadvantage—the advantage-building motive—and to prevent its rivals from acquiring competitive advantage—the advantage-denying motive. Network effects, in particular, are a mirror image of learning-by-doing in the sense that economies of scale originate in demand rather than cost. We therefore expect our analysis to extend beyond our specific setting.²

²A caveat is that consumers in our learning-by-doing model are short-lived. We leave the analysis of dynamic competition with forward-looking consumers to future research.

While the literature has generally focused on characterizing equilibrium behavior rather than anatomizing its implications for welfare, it has raised some red flags regarding the efficiency of dynamic competition. Bergemann and Välimäki (1997) show that the equilibrium in their experience-goods model is inefficient because of an informational externality between firms (see also Rob 1991 and Bolton and Harris 1999). The literature on network effects shows that an inferior product can be adopted and persist as the standard either because consumers' expectations are misaligned or favor the inferior product (Biglaiser and Crémer 2016; Halaburda, Jullien, and Yehezkel 2016). Nevertheless, none of these papers directly addresses the question whether the investment role of price is socially beneficial.

The remainder of this paper is organized as follows. Section I sets up our learning-by-doing model. Section II develops the first-best planner problem and the welfare metrics we use in the subsequent analysis. Sections III and IV present our numerical analysis of equilibria and associated deadweight loss over a wide range of the parameter space and summarize the regularities that emerge. Section V introduces our decomposition of deadweight loss and provides analytical bounds on some of its components. Section VI summarizes and concludes. An online Appendix contains additional details and results.

Throughout the paper we distinguish between propositions, whose proofs are in the Appendix, and results that summarize our exploration of the parameter space by the percentage of parameterizations that display an outcome of interest. We think of these percentages as characterizing frequencies of occurrence over the economically interesting portion of parameter space for two reasons. First, we proceed as if randomly sampling parameterizations from this portion of parameter space. Second, we either explore the entire feasible range of a parameter or, if lacking a bound, explore the parameter to a point beyond which further changes in its value do not change the equilibrium very much (see Section IIIA). While further extending the region of exploration may affect the reported percentages, this would be an artifact of repeatedly sampling basically identical equilibria.

I. Model

We use the discrete-time, infinite-horizon dynamic stochastic game between two firms in an industry characterized by learning-by-doing in Besanko, Doraszelski, and Kryukov (2014). We briefly review the model and refer the reader to Besanko, Doraszelski, and Kryukov (2014) and the online Appendix for details.

At any point in time, firm $n \in \{1, 2\}$ is described by its state $e_n \in \{0, 1, \dots, M\}$. State $e_n = 0$ indicates a potential entrant. State $e_n \in \{1, \dots, M\}$ indicates the cumulative experience or stock of know-how of an incumbent firm. By making a sale in the current period, an incumbent firm adds to its stock of know-how and, through learning-by-doing, lowers its production cost in the subsequent period. Competitive advantage and industry leadership are therefore determined endogenously in the model. If $e_1 > e_2$ ($e_1 < e_2$), then we refer to firm 1 as the leader (follower) and to firm 2 as the follower (leader). The industry's state is the vector of firms' states $\mathbf{e} = (e_1, e_2)$.

In each period, firms first set prices and then decide on exit and entry.³ During the price-setting phase, the state changes from \mathbf{e} to \mathbf{e}' . If firm 1 makes the sale, the state changes to $\mathbf{e}' = \mathbf{e}^{1+} = (\min\{e_1 + 1, M\}, e_2)$; if firm 2 makes the sale, the state changes to $\mathbf{e}' = \mathbf{e}^{2+} = (e_1, \min\{e_2 + 1, M\})$.

During the exit-entry phase, the state then changes from \mathbf{e}' to \mathbf{e}'' . Entry of firm n is a transition from state $e'_n = 0$ to state $e''_n = 1$ and exit is a transition from state $e'_n \geq 1$ to state $e''_n = 0$. As the exit of an incumbent firm creates an opportunity for a potential entrant to enter the industry, reentry is possible. The state at the end of the current period finally becomes the state at the beginning of the subsequent period.

Learning-by-Doing and Production Cost.—Incumbent firm n 's marginal cost $c(e_n) = \kappa\rho^{\log_2(\min\{e_n, m\})}$ depends on its stock of know-how e_n through a learning curve with a progress ratio $\rho \in [0, 1]$. A lower progress ratio implies stronger learning economies. A firm without prior experience has marginal cost $c(1) = \kappa > 0$. Once the firm reaches state m , there are no further experience-based cost reductions. We refer to an industry in state \mathbf{e} as a mature duopoly if $\mathbf{e} \geq (m, m)$ and as a mature monopoly if either $e_1 \geq m$ and $e_2 = 0$ or $e_1 = 0$ and $e_2 \geq m$.

Demand and Consumer Surplus.—One buyer enters the market each period and purchases one unit of either one of the “inside goods” offered by the incumbent firms at prices $\mathbf{p} = (p_1, p_2)$ or a competitively supplied “outside good” at a price p_0 equal to its marginal cost $c_0 \geq 0$. The probability that firm n makes the sale is given by the logit specification

$$D_n(\mathbf{p}) = \frac{\exp\left(\frac{v - p_n}{\sigma}\right)}{\sum_{k=0}^2 \exp\left(\frac{v - p_k}{\sigma}\right)} = \frac{\exp\left(\frac{-p_n}{\sigma}\right)}{\sum_{k=0}^2 \exp\left(\frac{-p_k}{\sigma}\right)},$$

where v is gross utility and $\sigma > 0$ is a scale parameter that governs the degree of product differentiation. As $\sigma \rightarrow 0$, goods become homogeneous. The price of the outside good $p_0 = c_0$ determines the overall level of demand for the inside goods. As it decreases, the market becomes smaller.

For future reference, the consumer surplus associated with our logit specification is

$$(1) \quad CS(\mathbf{p}) = \sigma \ln \left(\sum_{n=0}^2 \exp\left(\frac{v - p_n}{\sigma}\right) \right) = v + \sigma \ln \left(\sum_{n=0}^2 \exp\left(\frac{-p_n}{\sigma}\right) \right).$$

Scrap Values and Setup Costs.—If incumbent firm n exits the industry, it receives a scrap value X_n drawn from a symmetric triangular distribution $F_X(\cdot)$ with support $[\bar{X} - \Delta_X, \bar{X} + \Delta_X]$, where $E_X(X_n) = \bar{X}$ and $\Delta_X > 0$ is a scale parameter.⁴ If potential entrant n enters the industry, it incurs a setup cost S_n drawn from a symmetric triangular distribution $F_S(\cdot)$ with support $[\bar{S} - \Delta_S, \bar{S} + \Delta_S]$, where

³In the online Appendix, we show that our model is invariant to reversing this order.

⁴In the online Appendix, we show that our model is equivalent to a model with per-period, avoidable fixed costs but without scrap values.

$E_S(S_n) = \bar{S}$ and $\Delta_S > 0$ is a scale parameter. Scrap values and setup costs are independently and identically distributed across firms and periods, and their realization is observed by the firm but not its rival. Consequently, its rival perceives the firm as if it was following a mixed strategy.

A. Firms' Decisions

Firms aim to maximize the expected net present value (NPV) of their future cash flows, using discount factor $\beta \in [0, 1)$.⁵ We analyze their decisions by working backward from the exit-entry phase to the price-setting phase.

Exit and Entry Decisions.—To simplify the exposition, we focus on firm 1 and let $\phi_1(\mathbf{e}')$ denote the probability that firm 1 decides *not* to operate in state \mathbf{e}' : if $e_1 > 0$ so that firm 1 is an incumbent, then $\phi_1(\mathbf{e}')$ is the probability of exiting; if $e_1 = 0$ so that firm 1 is an entrant, then $\phi_1(\mathbf{e}')$ is the probability of not entering.

Pricing Decision.—The pricing decision $p_1(\mathbf{e})$ of incumbent firm 1 in state \mathbf{e} is uniquely determined (given $p_2(\mathbf{e})$) by the first-order condition

$$(2) \quad p_1(\mathbf{e}) - \frac{\sigma}{1 - D_1(\mathbf{p}(\mathbf{e}))} - c(e_1) + [U_1(\mathbf{e}^{1+}) - U_1(\mathbf{e})] \\ + \Upsilon(p_2(\mathbf{e})) [U_1(\mathbf{e}) - U_1(\mathbf{e}^{2+})] = 0,$$

where $\mathbf{p}(\mathbf{e}) = (p_1(\mathbf{e}), p_2(\mathbf{e}))$, $\Upsilon(p_2(\mathbf{e})) = \frac{D_2(\mathbf{p}(\mathbf{e}))}{1 - D_1(\mathbf{p}(\mathbf{e}))}$ is the probability of firm 2 making a sale conditional on firm 1 not making a sale, and $U_1(\mathbf{e}')$ denotes the continuation value of firm 1 in state \mathbf{e}' going into the exit-entry phase.

The pricing decision impounds two distinct goals beyond static profit $D_1(\mathbf{p}(\mathbf{e}))(p_1(\mathbf{e}) - c(e_1))$. The *advantage-building motive* $U_1(\mathbf{e}^{1+}) - U_1(\mathbf{e})$ is the reward the firm receives by *winning* a sale and moving down its learning curve. The *advantage-denying motive* $U_1(\mathbf{e}) - U_1(\mathbf{e}^{2+})$ is the penalty the firm avoids by *preventing its rival from winning* the sale and moving down its learning curve. The advantage-building and advantage-denying motives arise in a broad class of dynamic models and together capture the investment role of price.

⁵Besides time preference, the discount factor accounts for the probability that the industry “dies” for exogenous reasons (Besanko et al. 2010). Especially in technology industries, firms’ cost advantages may be rendered worthless by rapid product innovation from outside the industry that displaces their products. We use the same discount factor for firms and the first-best planner to avoid a “money pump” in which the government loans firms an arbitrarily large amount of money to make both better off. We discuss diverging discount factors in the online Appendix.

B. Equilibrium and Industry Dynamics

Because the demand and cost specification is symmetric, we restrict ourselves to symmetric Markov perfect equilibria (MPE) in pure strategies.⁶ Existence follows from the arguments in Doraszelski and Satterthwaite (2010). In a symmetric equilibrium, the decisions taken by firm 2 in state \mathbf{e} are identical to the decisions taken by firm 1 in state (e_2, e_1) .

Despite the restriction to symmetric equilibria, there is a substantial amount of multiplicity. Because the literature offers little guidance regarding equilibrium selection, we view all equilibria that arise for a fixed set of primitives as equally likely. In doing so, we take the ex ante perspective of a regulator that over time will confront a sequence of industries but at present has little detailed knowledge about these industries.

To study the evolution of the industry under a particular equilibrium, we use the policy functions \mathbf{p}_1 and ϕ_1 with typical element $p_1(\mathbf{e})$, respectively, $\phi_1(\mathbf{e}')$ to construct the matrix of state-to-state transition probabilities that characterizes the Markov process of industry dynamics. From this, we compute the transient distribution over states in period t , μ_t , starting from state $(0, 0)$ (the empty industry with just the outside good) in period 0. In line with our ex ante perspective, we have in mind a nascent industry in which two firms have developed new products that can potentially draw customers away from an established product (the outside good) but have not yet brought them to market.⁷ The typical element $\mu_t(\mathbf{e})$ is the probability that the industry is in state \mathbf{e} in period t . Depending on t , the transient distributions can capture short-run or long-run (steady-state) dynamics. We think of period 500 as the long run and, with a slight abuse of notation, denote μ_{500} by μ_∞ . We do not use the limiting (or ergodic) distribution to capture long-run dynamics because the Markov process implied by the equilibrium may have multiple closed communicating classes.

II. First-Best Planner, Welfare, and Deadweight Loss

A. First-Best Planner

Our welfare benchmark is a first-best planner who makes pricing, exit, and entry decisions to maximize the expected NPV of total surplus (consumer plus producer surplus). In contrast to the market, the planner centralizes and coordinates decisions across firms as in Bolton and Farrell (1990). To stack the deck against finding small deadweight losses, we assume an omniscient planner that has access to privately known scrap values and setup costs.

⁶Treating firms' decisions as decentralized and uncoordinated by focusing on symmetric equilibria stacks the deck against finding small deadweight losses. While asymmetric or correlated equilibria can avoid wasteful duplication and delay arising from entry and exit, as Bolton and Farrell (1990) discuss it is far from clear how the firms will "find" one of these equilibria without some process of explicit coordination.

⁷Starting from state $(0, 0)$ stacks the deck against finding small deadweight losses by fully recognizing any distortions in the entry process. It is also an interesting setting in its own right as sellers of next-generation products aiming to establish a "footprint" are a pervasive feature of the business landscape, and one where the investment role of price is particularly salient.

We refer the reader to the online Appendix for a formal statement of the first-best planner problem. From the contraction mapping theorem, a solution exists and is unique. We use it to construct the matrix of state-to-state transition probabilities and compute the transient distribution over states in period t , μ_t^{FB} , starting from state $(0, 0)$ in period 0.

B. Welfare and Deadweight Loss

To capture both short-run and long-run dynamics, our welfare metric is the expected NPV of total surplus

$$(3) \quad TS_\beta = \sum_{t=0}^{\infty} \beta^t \sum_{\mathbf{e}} \mu_t(\mathbf{e}) TS(\mathbf{e}),$$

where $TS(\mathbf{e}) = CS(\mathbf{e}) + PS(\mathbf{e})$ is total surplus in state \mathbf{e} under a particular equilibrium, $CS(\mathbf{e}) = CS(\mathbf{p}(\mathbf{e}))$ is given by equation (1), and $PS(\mathbf{e})$ includes the static profit $\Pi(\mathbf{e}) = \sum_{n=1}^2 D_n(\mathbf{p}(\mathbf{e}))(p_n(\mathbf{e}) - c(e_n))$ of incumbent firms as well as their expected scrap values and the expected setup costs of potential entrants (see Appendix A).

Under the first-best planner solution, we define the expected NPV of total surplus TS_β^{FB} analogously. The deadweight loss arising in equilibrium therefore is

$$(4) \quad DWL_\beta = TS_\beta^{FB} - TS_\beta.$$

Because DWL_β is measured in arbitrary monetary units, we normalize it by the *maximum value added* by the industry $VA_\beta = TS_\beta^{FB} - TS_\beta^\emptyset$, where $TS_\beta^\emptyset = \frac{v - p_0}{1 - \beta}$ is the expected NPV of total surplus if the industry remains empty forever with just the outside good. Note that VA_β is an upper bound on the contribution of the inside goods to the expected NPV of total surplus. We refer to DWL_β/VA_β as the *relative deadweight loss*.⁸

III. Numerical Analysis and Equilibrium

In the online Appendix, we show in an analytically tractable special case of our model with a two-step learning curve ($M = m = 2$), homogeneous goods ($\sigma = 0$), and mixed exit and entry strategies ($\Delta_X = \Delta_S = 0$) that even if pricing is efficient, exit and entry may not be. Distortions in exit and entry can take the form of over-exit (too much or early exit), under-exit (too little or late exit), over-entry (too much or early entry), under-entry (too little or late entry), and cost-inefficient exit where the lower-cost firm exits the industry while the higher-cost firm does not.

⁸We normalize by VA_β because, like DWL_β but unlike TS_β^{FB} and TS_β , it does not depend on gross utility v . While v does not affect the behavior of industry participants in any way, we set it to make the expected NPV of consumer surplus CS_β positive. An alternative is to set $v = p_0$, which implies $VA_\beta = TS_\beta^{FB}$.

While theoretical analysis enables us to establish that dynamic competition is not necessarily fully efficient, it is ill-suited to answer the question of how efficient dynamic competition is. We therefore turn to numerical analysis.

A. Parameterization and Computation

Our learning-by-doing model has four key parameters: the progress ratio ρ , the degree of product differentiation σ , the price of the outside good p_0 , and the expected scrap value \bar{X} . Robustness checks indicate that the remaining parameters play a lesser role or matter mostly in relation to the four key parameters (Besanko et al. 2010; Besanko, Doraszelski, and Kryukov 2014).

To explore how the equilibria vary with the parameters and to search for multiple equilibria in a systematic fashion, we use the homotopy method to compute six two-dimensional slices through the equilibrium correspondence along (ρ, σ) , (ρ, p_0) , (ρ, \bar{X}) , (σ, p_0) , (σ, \bar{X}) , and (p_0, \bar{X}) . We choose sufficiently large upper bounds for σ and p_0 so that beyond them “things don’t change much.” Back-of-the-envelope calculations yield $\sigma \leq 10$ and $p_0 \leq 20$. We exclude equilibria for which the industry is not viable in the sense that the probability $1 - \phi_1(0, 0)^2$ that the industry “takes off” is below 0.01. Throughout we hold the remaining parameters fixed at the values in Table 1. While this baseline parameterization is not intended to be representative of any particular industry, it is neither entirely unrepresentative nor extreme.

Due to the large number of parameterizations and multiplicity of equilibria, we require a way to summarize them. In a first step, we average an outcome of interest over the equilibria at a parameterization. This “random sampling” is in line with our decision to refrain from equilibrium selection and ensures that parameterizations with many equilibria carry the same weight as parameterizations with few equilibria.

In a second step, we randomly sample parameterizations. To make this practical, we represent a two-dimensional slice through the equilibrium correspondence with a grid of values for the parameters spanning the slice. Table 1 lists the grid points we use for the four key parameters. We mostly use uniformly spaced grid points, except for $\sigma > 1$, where the grid points approximate a log scale in order to explore very high degrees of product differentiation. We associate each point in a two-dimensional grid with the corresponding average over equilibria. We then pool the points on the six slices through the equilibrium correspondence and obtain the distribution of the outcome of interest.

B. Equilibrium and First-Best Planner Solution

To illustrate the types of behavior that can emerge in our learning-by-doing model, Table 2 shows several metrics of industry structure, conduct, and performance for two of the three equilibria that arise at the baseline parameterization in Table 1 as well as for the first-best planner solution.⁹

⁹We refer the reader to the online Appendix for formal definitions of these metrics. The third equilibrium is essentially intermediate between the two shown in Table 2.

TABLE 1—BASELINE PARAMETERIZATION AND GRID POINTS

Parameter	Value	Grid
Maximum stock of know-how M	30	
Cost at top of learning curve κ	10	
Bottom of learning curve m	15	
Progress ratio ρ	0.75	$\rho \in \{0, 0.05, \dots, 1\}$
Gross utility v	10	
Product differentiation σ	1	$\sigma \in \{0.2, 0.3, \dots, 1, 1.3, 1.6, 2, 2.5, 3, 2, 4, 5, 6.3, 7.9, 10\}$
Price of outside good p_0	10	$p_0 \in \{0, 1, \dots, 20\}$
Scrap value \bar{X}, Δ_X	1.5, 1.5	$\bar{X} \in \{-1.5, -1, \dots, 7.5\}$
Setup cost \bar{S}, Δ_S	4.5, 1.5	
Discount factor β	0.9524	

TABLE 2—INDUSTRY STRUCTURE, CONDUCT, AND PERFORMANCE

	Aggressive equilibrium	Accommodative equilibrium	Planner solution
<i>Structure</i>			
Expected short-run number of firms N_1	1.92	1.91	1.00
Expected long-run number of firms N_∞	1.08	2.00	1.00
<i>Conduct</i>			
Expected short-run average price \bar{p}_1	-31.29	5.32	4.28
Expected long-run average price \bar{p}_∞	8.28	5.24	3.25
Expected time to maturity T^m	19.09	37.54	15.02
<i>Performance</i>			
Expected NPV of consumer surplus CS_β	93.87	103.29	131.66
Expected NPV of total surplus TS_β	96.02	105.45	110.45
Deadweight loss DWL_β	14.43	5.01	—
Relative deadweight loss DWL_β/VA_β	13.06%	4.54%	—

Notes: Aggressive and accommodative equilibrium, and first-best planner solution. Baseline parameterization.

The first-best planner operates the industry as a natural monopoly.¹⁰ In state $(0, 0)$ in period 0, the planner decides to operate a single firm (say firm 1) in the subsequent period; the expected short-run number of firms in Table 2 is thus $N_1^{FB} = 1.00$. In period $t \geq 1$, the planner marches firm 1 down its learning curve. We define the expected time to maturity as the expected time until the industry first becomes either a mature monopoly or a mature duopoly. The value $T^{m,FB} = 15.02$ measures the speed at which the planner moves firm 1 down its learning curve. At the bottom of its learning curve firm 1 charges a price equal to marginal cost; the expected long-run average price is thus $\bar{p}_\infty^{FB} = 3.25$.

Both equilibria differ considerably from the first-best planner solution. Behavior in what we call an *aggressive equilibrium* in Besanko, Doraszelski, and Kryukov (2014) resembles conventional notions of predatory pricing.¹¹ The pricing decision exhibits a deep well in state $(1, 1)$ with $p_1(1, 1) = -34.78$; this is reflected in the

¹⁰Outside the baseline parameterization in Table 1, the first-best planner does not necessarily operate the industry as a natural monopoly. For example, if the degree of product differentiation is sufficiently large, then the planner immediately decides to operate both firms and continues to do so as they move down their learning curves.

¹¹In Besanko, Doraszelski, and Kryukov (2014) we formalize the notion of predatory pricing and disentangle it from mere competition for efficiency on a learning curve.

expected short-run average price $\bar{p}_1 = -31.29$ in Table 2. A *well* is a preemption battle where firms vie to be the first to move down from the top of their learning curves. Such a battle is likely to ensue because, as indicated by $N_1 = 1.92$, both firms are likely to enter the industry in period 0. After the industry has emerged from the preemption battle in state $(1, 1)$, the leader (say firm 1) continues to price aggressively. Indeed, the pricing decision exhibits a deep trench along the e_1 -axis with $p_1(e_1, 1)$ ranging from 0.08 to 1.24 for $e_1 \in \{2, \dots, 30\}$. A *trench* is a price war that the leader wages against the follower. In the trench the follower (firm 2) exits the industry with a positive probability of $\phi_2(e_1, 1) = 0.22$ for $e_1 \in \{2, \dots, 30\}$. The follower remains in this exit zone as long as it does not win a sale. Once the follower exits, the leader raises its price and the industry becomes an entrenched monopoly. The industry may also evolve into a mature duopoly if the follower manages to win a sale but this is far less likely than an entrenched monopoly; indeed, the expected long-run number of firms is $N_\infty = 1.08$.

In the *accommodative equilibrium*, by contrast, the pricing decision exhibits a shallow well in state $(1, 1)$ with $p_1(1, 1) = 5.05$; in Table 2 this is reflected in $\bar{p}_1 = 5.32$. A much milder preemption battle is again likely to ensue, as indicated by $N_1 = 1.91$. After the industry has emerged from the preemption battle in state $(1, 1)$, the leader enjoys a competitive advantage over the follower. Without mobility barriers in the form of trenches, however, this advantage is temporary and, as indicated by $N_\infty = 2.00$, the industry evolves into a mature duopoly. While the aggressive equilibrium mainly involves over-entry, the accommodative equilibrium therefore involves both over-entry and under-exit. Learning economies are exhausted slowest in the accommodative equilibrium with $T^m = 37.54$, followed by the aggressive equilibrium with $T^m = 19.09$ and the first-best planner solution with $T^{m,FB} = 15.02$. This large gap arises because sales are split between the inside goods in the accommodative equilibrium, as well as at least initially with the outside good.

As the industry is substantially more likely to be monopolized in the aggressive equilibrium than in the accommodative equilibrium, in Table 2 the expected long-run average price is $\bar{p}_\infty = 8.28$ compared to $\bar{p}_\infty = 5.24$. The expected NPV of consumer surplus CS_β is lower, as is the expected NPV of total surplus TS_β . Consequently, the deadweight loss DWL_β is higher in the aggressive equilibrium than in the accommodative equilibrium. However, the relative deadweight loss DWL_β/VA_β seems modest in both equilibria: 13.06 percent of the maximum value added by the industry in the aggressive equilibrium and 4.54 percent in the accommodative equilibrium.

IV. Does Dynamic Competition Lead to Low Deadweight Loss?

The relative deadweight loss DWL_β/VA_β is modest more generally. Summarizing a large number of parameterizations and equilibria, Figure 1 shows the cumulative distribution function (CDF) of DWL_β/VA_β as a solid line. Result 1 highlights some findings.

RESULT 1: *The relative deadweight loss DWL_β/VA_β is less than 0.05, 0.1, and 0.2 in 26.40 percent, 65.83 percent, and 92.03 percent of parameterizations, respectively. The median of DWL_β/VA_β is 0.0777.*




FIGURE 1. DISTRIBUTION OF RELATIVE DEADWEIGHT LOSS DWL_{β}/VA_{β}

Notes: All equilibria (solid line), best equilibrium (dotted line), and worst equilibrium (dashed line). Parameterizations and equilibria within parameterizations weighted equally.

There is a large relative deadweight loss DWL_{β}/VA_{β} for a small number of parameterizations. Under these parameterizations, the industry almost fails to take off $(1 - \phi_1(0,0))^2 \approx 0.01$ because the outside good is highly attractive. Near this “cusp of viability” the contribution of the inside goods to the expected NPV of total surplus is small and thus $VA_{\beta} \approx 0$.

Recall that we average over equilibria at a given parameterization to obtain the distribution of DWL_{β}/VA_{β} . To look behind these averages, we consider the *best equilibrium* with the highest value of TS_{β} at a given parameterization as well as the *worst equilibrium* with the lowest value of TS_{β} . Figure 1 shows the resulting distributions of DWL_{β}/VA_{β} using a dotted line for the best equilibrium and a dashed line for the worst equilibrium.

RESULT 2: (i) For the best equilibrium, the relative deadweight loss DWL_{β}/VA_{β} is less than 0.05, 0.1, and 0.2 in 44.25 percent, 71.11 percent, and 92.10 percent of parameterizations, respectively. The median of DWL_{β}/VA_{β} is 0.0571. (ii) For the worst equilibrium, the relative deadweight loss DWL_{β}/VA_{β} is less than 0.05, 0.1, and 0.2 in 18.67 percent, 56.40 percent, and 91.80 percent of parameterizations, respectively. The median of DWL_{β}/VA_{β} is 0.0922.

Hence, even in the worst equilibria the relative deadweight loss DWL_{β}/VA_{β} is modest for a wide range of parameterizations.

In Appendix B we offer formal definitions of aggressive and accommodative equilibria and link them to the worst, respectively, best equilibria. Although this link is not perfect, to build intuition, in what follows we associate the best equilibrium with an accommodative equilibrium and the worst equilibrium, to the extent that it differs from the best equilibrium, with an aggressive equilibrium. If the equilibrium is unique, then we associate it with an accommodative equilibrium.

A. Deadweight Loss in Perspective: Static Non-Cooperative Pricing Counterfactual and Collusive Solution

Is a relative deadweight loss DWL_{β}/VA_{β} of 10 percent “small” and one of 30 percent “large?” To put these percentages in perspective, we first show that the deadweight loss is lower than expected in light of the traditional view of price as affecting current profit but not the evolution of the industry. To this end, we shut down the investment role of price, leaving incumbent firms to maximize static profit in the price-setting phase. We refer the reader to the online Appendix for a formal statement of the static non-cooperative pricing counterfactual.

The investment role of price is by and large socially beneficial. Using a solid line, Figure 2 shows the distribution of the deadweight loss ratio $DWL_{\beta}^{SN}/DWL_{\beta}$ that compares the static non-cooperative pricing counterfactual to the equilibrium and Result 3 summarizes.

RESULT 3: DWL_{β}^{SN} is at least as large as DWL_{β} in 80.32 percent of parameterizations and at least twice as large in 45.45 percent of parameterizations. The median of $DWL_{\beta}^{SN}/DWL_{\beta}$ is 1.8328.

In a second benchmark for the deadweight loss that arises in equilibrium, we assume that firms collude by centralizing and coordinating their pricing, exit, and entry decisions to maximize the expected NPV of producer surplus. We again refer the reader to the online Appendix. Figure 2 shows the distribution of the deadweight loss ratio $DWL_{\beta}^{CO}/DWL_{\beta}$ using a dashed line and Result 4 summarizes.

RESULT 4: DWL_{β}^{CO} is at least as large as DWL_{β} in 94.74 percent of parameterizations and at least twice as large in 36.32 percent of parameterizations. The median of $DWL_{\beta}^{CO}/DWL_{\beta}$ is 1.4701.

In sum, the static non-cooperative pricing counterfactual and the collusive solution reinforce that dynamic competition generally leads to low deadweight loss. They also show that a low deadweight loss is almost certainly not hardwired into the primitives of our learning-by-doing model. Instead, there is something in the investment role of price and the nature of dynamic competition that in equilibrium leads to low deadweight loss.

B. Differences between Equilibria and First-Best Planner Solution

There are typically substantial differences between the equilibria and the first-best planner solution. As we show in the online Appendix, the equilibria typically have too many firms in the short run, consistent with over-entry, and too many firms in the long run, consistent with under-exit. This latter tendency is exacerbated in the best equilibrium. The speed of learning in the equilibria is generally too slow. The best equilibrium exhausts learning economies even more slowly than the worst equilibrium because pricing is initially less aggressive and more firms split sales in an accommodative equilibrium than in an aggressive equilibrium.

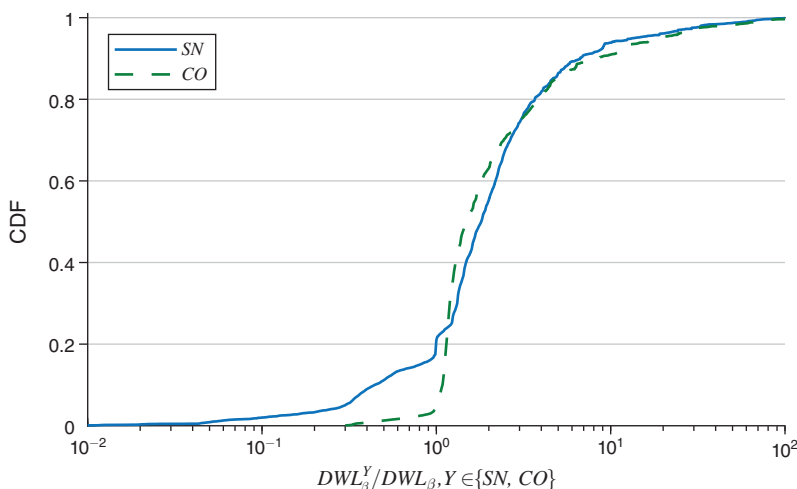


FIGURE 2. DISTRIBUTION OF DEADWEIGHT LOSS RATIO $DWL_{\beta}^{SN}/DWL_{\beta}$ AND $DWL_{\beta}^{CO}/DWL_{\beta}$

Notes: Static non-cooperative pricing counterfactual (solid line) and collusive solution (dashed line). Log scale. Parameterizations and equilibria within parameterizations weighted equally.

V. Why Does Dynamic Competition Lead to Low Deadweight Loss?

Section IV leaves us with a puzzle. Dynamic competition leads to low deadweight loss, but this is not because the equilibrium resembles the first-best planner solution. On the contrary, the best equilibrium can differ even more from the planner solution than the worst equilibrium.

A. Decomposition

The deadweight loss DWL_{β} in equation (4) comprises differences between the equilibrium and the first-best planner solution in any given state and in the states that are visited over time. We accordingly decompose it as $DWL_{\beta} = DWL_{\beta}^{PR} + DWL_{\beta}^{EE} + DWL_{\beta}^{MS}$:

$$\begin{aligned}
 DWL_{\beta}^{PR} &= \sum_{t=0}^{\infty} \beta^t \sum_{\mathbf{e}} \mu_t(\mathbf{e}) [CS^{FB}(\mathbf{e}) + \Pi^{FB}(\mathbf{e}) - (CS(\mathbf{e}) + \Pi(\mathbf{e}))], \\
 DWL_{\beta}^{EE} &= \sum_{t=0}^{\infty} \beta^t \sum_{\mathbf{e}} \mu_t(\mathbf{e}) [PS^{FB}(\mathbf{e}) - \Pi^{FB}(\mathbf{e}) - (PS(\mathbf{e}) - \Pi(\mathbf{e}))], \\
 DWL_{\beta}^{MS} &= \sum_{t=0}^{\infty} \beta^t \sum_{\mathbf{e}} [\mu_t^{FB}(\mathbf{e}) - \mu_t(\mathbf{e})] TS^{FB}(\mathbf{e}),
 \end{aligned}$$

where $CS(\mathbf{e}) + \Pi(\mathbf{e})$ is the sum of consumer surplus and static profit in state \mathbf{e} and $PS(\mathbf{e}) - \Pi(\mathbf{e})$ is the difference of producer surplus and static profit and thus is the part of producer surplus that accounts for scrap values and setup costs. We analogously define $CS^{FB}(\mathbf{e}) + \Pi^{FB}(\mathbf{e})$ and $PS^{FB}(\mathbf{e}) - \Pi^{FB}(\mathbf{e})$ for the first-best planner solution.

The *pricing distortion* DWL_{β}^{PR} is the deadweight loss from state-wise differences in pricing conduct between the equilibrium and the first-best planner solution. The *exit and entry distortion* DWL_{β}^{EE} is similarly the deadweight loss from state-wise differences in exit and entry conduct. Expected inflows from scrap values in state \mathbf{e} contribute positively to $PS(\mathbf{e}) - \Pi(\mathbf{e})$ and expected outflows from setup costs negatively. A positive value of DWL_{β}^{EE} thus reflects a tendency for over-entry or under-exit relative to the planner solution while a negative value reflects a tendency for under-entry or over-exit.

The *market structure distortion* DWL_{β}^{MS} is the deadweight loss from differences in the evolution of the industry. A negative value of DWL_{β}^{MS} indicates that the equilibrium puts more weight on more favorable market structures with higher values of $TS^{FB}(\mathbf{e})$ than the first-best planner solution; a positive value indicates the reverse. The state \mathbf{e} completely describes the number of incumbent firms, and therefore the extent of product variety, along with their cost positions. The value of $TS^{FB}(\mathbf{e})$ is high if firms' cost positions in relation to the price of the outside good yield large gains from trade; $TS^{FB}(\mathbf{e})$ is also high if a large number of firms fosters product variety or allows the planner to receive scrap values by ceasing to operate excess firms. The factors contributing to a negative value of DWL_{β}^{MS} are therefore over-entry and under-exit as well as fast exploitation of learning economies. The factors contributing to a positive value of DWL_{β}^{MS} are under-entry, over-exit, slow exploitation of learning economies, and cost-inefficient exit. Cost-inefficient exit contributes to a positive value of DWL_{β}^{MS} in the same way as slow exploitation of learning economies by rendering firms' cost positions less favorable.

Because DWL_{β}^{EE} and DWL_{β}^{MS} depend on over-entry and under-exit in opposite ways, they can offset each other. Accordingly we define the *non-pricing distortion* as $DWL_{\beta}^{NPR} = DWL_{\beta}^{EE} + DWL_{\beta}^{MS}$. It reflects (i) the *net* social loss from a suboptimal number of firms (setup costs net of scrap values net of social benefits of product variety), (ii) the *gross* social loss from a suboptimal exploitation of learning economies, and (iii) the *gross* social loss from cost-inefficient exit.

Figure 3 shows the distribution of $DWL_{\beta}^{PR}/DWL_{\beta}$, $DWL_{\beta}^{EE}/DWL_{\beta}$, $DWL_{\beta}^{MS}/DWL_{\beta}$, and $DWL_{\beta}^{NPR}/DWL_{\beta}$ for a large number of parameterizations and equilibria. We scale each term of the decomposition by DWL_{β} to better gauge its size. Result 5 highlights some findings.

RESULT 5: (i) The pricing distortion DWL_{β}^{PR} is positive in 96.44 percent of parameterizations.¹² The median of $DWL_{\beta}^{PR}/DWL_{\beta}$ is 0.6565. (ii) The exit and entry distortion DWL_{β}^{EE} is positive in 81.26 percent of parameterizations. The median of $DWL_{\beta}^{EE}/DWL_{\beta}$ is 0.5728. (iii) The market structure distortion DWL_{β}^{MS} is negative in 70.07 percent of parameterizations. The median of $DWL_{\beta}^{MS}/DWL_{\beta}$ is -0.2279 . (iv) The non-pricing distortion DWL_{β}^{NPR} is positive in 92.03 percent of parameterizations. The median of $DWL_{\beta}^{NPR}/DWL_{\beta}$ is 0.3435.

¹²We take DWL_{β}^{PR} to be zero if $|DWL_{\beta}^{PR}/DWL_{\beta}| < 0.001$. We proceed analogously for the other terms.




FIGURE 3. DISTRIBUTION OF $DWL_{\beta}^{PR}/DWL_{\beta}$, $DWL_{\beta}^{EE}/DWL_{\beta}$, $DWL_{\beta}^{MS}/DWL_{\beta}$, AND $DWL_{\beta}^{NPR}/DWL_{\beta}$

Note: Parameterizations and equilibria within parameterizations weighted equally.

Note that DWL_{β}^{EE} is typically positive by part (ii) of Result 5 because of over-entry and under-exit.¹³ Also because of over-entry and under-exit, DWL_{β}^{MS} is typically negative by part (iii).¹⁴ Thus, as highlighted in Result 6, DWL_{β}^{MS} typically offsets DWL_{β}^{EE} so that DWL_{β}^{NPR} is very often smaller than its largest component (in absolute value).

RESULT 6: $|DWL_{\beta}^{NPR}|$ is smaller than $\max\{|DWL_{\beta}^{EE}|, |DWL_{\beta}^{MS}|\}$ in 88.64 percent of parameterizations.

Parts (i) and (iv) of Result 5 show that generally both pricing and non-pricing distortions contribute to deadweight loss. As Result 7 shows, the pricing distortion is often larger than the non-pricing distortion.

RESULT 7: The pricing distortion DWL_{β}^{PR} is larger than the non-pricing distortion DWL_{β}^{NPR} in 73.23 percent of parameterizations.

Armed with these patterns in the decomposition terms, we delve into the mechanisms through which dynamic competition leads to low deadweight loss.

¹³ DWL_{β}^{EE} is zero in 14.87 percent of parameterizations and negative in just 3.87 percent of parameterizations (see again footnote 12).

¹⁴ DWL_{β}^{MS} can be positive if the degree of product differentiation σ is sufficiently large to ensure that the industry evolves into a mature duopoly under the equilibrium and the first-best planner solution. In this case, the positive value of DWL_{β}^{MS} reflects mainly the slow exploitation of learning economies.

B. Why Is the Best Equilibrium So Good?

The deadweight loss in the best equilibrium is small although it often differs greatly from the first-best planner solution. We show that this is because the pricing and non-pricing distortions are both small. Recall that we associate the best equilibrium with an accommodative equilibrium.

Why Is the Pricing Distortion Small?—At first blush, there appears to be no reason for the pricing distortion in an accommodative equilibrium to be particularly small. At the baseline parameterization, for example, equilibrium prices $(p_1(\mathbf{e}), p_2(\mathbf{e}))$ substantially exceed static first-best prices $(c_1(e_1), c_2(e_2))$ even once the industry becomes a mature duopoly.

Proposition 1, however, bounds the contribution $CS^{FB}(\mathbf{e}) + \Pi^{FB}(\mathbf{e}) - (CS(\mathbf{e}) + \Pi(\mathbf{e}))$ to the pricing distortion DWL_{β}^{PR} by the demand for the outside good $D_0(\mathbf{p}(\mathbf{e}))$.

PROPOSITION 1: *Consider a symmetric state $\mathbf{e} = (e, e)$, where $e > 0$. If $p_0 \geq \kappa$, $p_1(\mathbf{e}) > c(e)$, and $D_0(\mathbf{p}(\mathbf{e})) < 1/2$, then*

$$CS^{FB}(\mathbf{e}) + \Pi^{FB}(\mathbf{e}) - (CS(\mathbf{e}) + \Pi(\mathbf{e})) \leq \frac{(p_1(\mathbf{e}) - c(e))^2}{\sigma} D_0(\mathbf{p}(\mathbf{e})) (1 - D_0(\mathbf{p}(\mathbf{e}))).$$

Note that the bound in Proposition 1 approaches 0 as the demand for the outside good approaches 0. This often has bite: as the incumbent firms move down their learning curves, they improve their cost positions relative to the price of the outside good and drive the share of the outside good close to 0. Exceptions occur only if learning economies are weak— ρ is close to 1—or inconsequential because the outside good is highly attractive.

The bound in Proposition 1 transcends logit demand. To a first-order approximation, the deadweight loss due to market power decreases as demand becomes less price elastic (Harberger 1954). With logit demand, a decrease in the price elasticity of aggregate demand for the inside goods is associated with a decrease in the demand for the outside good. With linear demand, the aggregate demand for the inside goods also becomes less elastic as their prices, and with them the demand for the outside good, decrease (see online Appendix).

Why Is the Non-Pricing Distortion Small?—Recall that especially accommodative equilibria have too many firms, both in the short run and in the long run, consistent with over-entry and under-exit. Too many firms give rise to a social loss from wasteful duplication of setup costs due to over-entry and scrap values that firms forgo in equilibrium due to under-exit, but they also give rise to a social benefit from additional product variety. Negative values of DWL_{β}^{MS} thus tend to offset positive values of DWL_{β}^{EE} (Result 6) and the net social loss from too many firms is small. This is especially relevant because accommodative equilibria tend to arise when the degree of product differentiation σ is high (see Appendix B). Hence, the social benefit of additional product variety tends to be large.

Learning economies accentuate this “silver lining” from additional product variety. Suppose the industry evolves into a mature duopoly in an accommodative

equilibrium and into a mature monopoly in the first-best planner solution, and that there is no further exit and entry. Then, beyond a certain period, the equilibrium transient distribution $\mu_t(\mathbf{e})$ puts all mass on states \mathbf{e} , where $\mathbf{e} \geq (m, m)$, and the first-best transient distribution $\mu_t^{FB}(\mathbf{e})$ puts all mass on states \mathbf{e} , where either $e_1 \geq m$ and $e_2 = 0$ or $e_1 = 0$ and $e_2 \geq m$. Hence, in the market structure distortion DWL_{β}^{MS} , $\sum_{\mathbf{e}} [\mu_t^{FB}(\mathbf{e}) - \mu_t(\mathbf{e})] TS^{FB}(\mathbf{e}) = TS^{FB}(m, 0) - TS^{FB}(m, m) = CS^{FB}(m, 0) - CS^{FB}(m, m)$, where the second equality follows because the planner sets static first-best prices. Here, $CS^{FB}(m, 0) - CS^{FB}(m, m)$ can be thought of as the reduction in DWL_{β}^{MS} due to the social benefit of additional product variety. From equation (1) it is easy to show that

$$CS^{FB}(m, m) - CS^{FB}(m, 0) = \sigma \ln \left(\frac{\exp\left(\frac{-p_0}{\sigma}\right) + 2\exp\left(\frac{-c(m)}{\sigma}\right)}{\exp\left(\frac{-p_0}{\sigma}\right) + \exp\left(\frac{-c(m)}{\sigma}\right)} \right) < 0,$$

where $c(m) = \kappa\rho^{\log_2 m}$, and

$$\begin{aligned} & \frac{d(CS^{FB}(m, m) - CS^{FB}(m, 0))}{d\rho} \\ &= -\left(2D_1(c(m), c(m)) - D_1(c(m), \infty)\right) \kappa\rho^{\log_2 m - 1} \log_2 m < 0. \end{aligned}$$

Hence, as learning economies strengthen, the reduction in DWL_{β}^{MS} (and thus DWL_{β}^{NPR}) due to the social benefit of additional product variety increases.

C. Why Is the Worst Equilibrium Not So Bad?

Recall that we associate the worst equilibrium with an aggressive equilibrium. We show in two steps that the deadweight loss is small (albeit not as small as in the best equilibrium).

Why Is the Pricing Distortion Small?—There again appears to be no reason for the pricing distortion in an aggressive equilibrium to be small. At the baseline parameterization, the industry evolves into a mature monopoly. Proposition 2 bounds the contribution $CS^{FB}(\mathbf{e}) + \Pi^{FB}(\mathbf{e}) - (CS(\mathbf{e}) + \Pi(\mathbf{e}))$ to the pricing distortion DWL_{β}^{PR} by the degree of product differentiation σ and the advantage-building motive $U_1(\mathbf{e}^{1+}) - U_1(\mathbf{e})$.

PROPOSITION 2: *Consider a state $\mathbf{e} = (e, 0)$, where $e > 0$. Then*

$$CS^{FB}(\mathbf{e}) + \Pi^{FB}(\mathbf{e}) - (CS(\mathbf{e}) + \Pi(\mathbf{e})) < \begin{cases} \sigma & \text{if } 0 \leq U_1(\mathbf{e}^{1+}) - U_1(\mathbf{e}) < \sigma \left(1 + \exp\left(\frac{p_0 - c(e)}{\sigma}\right)\right) \\ \sigma + |U_1(\mathbf{e}^{1+}) - U_1(\mathbf{e})| & \text{otherwise.} \end{cases}$$

Note that the bound in Proposition 2 approaches σ as the incumbent firm moves down its learning curve and $U_1(\mathbf{e}^{1+}) - U_1(\mathbf{e})$ approaches 0.

Proposition 2 has bite because aggressive equilibria tend to arise when the degree of product differentiation σ is low (see Appendix B). This is intuitive: pricing aggressively to marginalize one's rival or altogether force it from the industry is especially attractive if products are close substitutes so that firms on an equal footing would fiercely compete.

Why Is the Non-Pricing Distortion Small?—While an aggressive equilibrium usually involves *delayed* exit (relative to the first-best planner solution), it involves rather brisk *eventual* exit. Thus, the industry usually quickly evolves toward the first-best market structure in an aggressive equilibrium, which tends to keep DWL_{β}^{MS} , and thus DWL_{β}^{NPR} , small. The small non-pricing distortion reflects the resulting relatively low net social loss from a suboptimal number of firms. Put another way, in an aggressive equilibrium, this net social loss is small because competition for the market resolves itself quickly and winnows out firms in a fairly efficient way.

VI. Summary and Conclusion

We study industries where price serves as an investment. The investment role arises when a firm's price affects not only its current profit but also its future competitive position vis-à-vis its rivals. Competition is dynamic as firms jostle for competitive advantage through the prices they set.

Our analysis suggests that dynamic competition is fairly efficient. Deadweight loss tends to be low (Results 1, 2, and 4) and the investment role of price socially beneficial (Result 3). This occurs even though equilibrium behavior and industry dynamics often differ markedly from the first-best planner solution.

So why is dynamic competition fairly efficient? The answer boils down to the key fundamental that gives rise to the investment role of price in the first place: learning-by-doing.

The pricing distortion tends to be the largest contributor to deadweight loss (part (i) of Result 5 and Result 7). Our bounds on the pricing distortion tighten as the incumbent firms move down their learning curves (Propositions 1 and 2). If the industry evolves into a mature duopoly in an accommodative equilibrium, this happens because the demand for the inside goods becomes less elastic as the incumbent firms improve their cost positions relative to the outside good and drive the share of the outside good close to zero. If the industry evolves into a mature monopoly in an aggressive equilibrium, the bound further tightens as the degree of product differentiation decreases and the threat of substitution to the outside good holds market power in check. Moreover, this is precisely when aggressive equilibria tend to arise in the first place.

The non-pricing distortion comprises the exit and entry distortion and the market structure distortion. In an aggressive equilibrium, the non-pricing distortion tends to be low despite the wasteful duplication of setup costs due to over-entry because competition for the market resolves itself quickly and winnows out firms in a fairly efficient way. In an accommodative equilibrium, the non-pricing distortion tends to be higher. However, the market structure distortion tends to partly offset the exit and entry distortion (parts (ii) and (iii) of Result 5 and Result 6). Too many firms give rise not only to a social loss from incurred setup costs due to over-entry and forgone

scrap values due to under-exit, but also to a social benefit from additional product variety. Learning economies accentuate this benefit by making product variety less costly to procure.

Put simply, when price serves as an investment, dynamic competition is fairly efficient because of the efficiency-enhancing properties of the investment.

From a regulatory perspective, our analysis suggests that in settings where price serves as an investment, the upside from interventions and policies aimed at shaping competition for the market is likely to be limited. As we show in Besanko, Doraszelski, and Kryukov (2014), blunt pricing conduct restrictions can lead to substantial welfare losses. Although more subtle pricing conduct restrictions can lead to welfare gains, typically by eliminating equilibria that entail predation-like behavior, achieving these gains in practice requires detailed knowledge of demand and cost primitives. Unless the regulator executes flawlessly based on this knowledge, it may therefore be preferable not to intervene at all and tolerate the “not so bad” welfare losses that typically arise under dynamic competition.

A caveat to our analysis is that interventions and policies remain appropriate that forestall collusion between firms that otherwise engage in competition for the market or exclusionary practices by these firms, such as forcing customers to sign long-term exclusive contracts. Our analysis should thus be interpreted as suggesting that there is little in the *fundamentals* of dynamic competition when price serves as an investment that suggests that antitrust actions have a sizable upside.

APPENDIX A. PRODUCER SURPLUS

Under decentralized exit and entry, producer surplus in state \mathbf{e} is $PS(\mathbf{e}) = \sum_{n=1}^2 PS_n(\mathbf{e})$, where producer surplus of firm 1 in state \mathbf{e} is

$$PS_1(\mathbf{e}) = D_1(\mathbf{p}(\mathbf{e}))(p_1(\mathbf{e}) - c(e_1)) \\ + \sum_{n=0}^2 D_n(\mathbf{p}(\mathbf{e})) \left\{ \mathbf{1}[e_1 > 0] \phi_1(\mathbf{e}^{n+}) E_X[X_1 | X_1 \geq \hat{X}_1(\mathbf{e}^{n+})] \right. \\ \left. - \mathbf{1}[e_1 = 0] (1 - \phi_1(\mathbf{e}^{n+})) E_S[S_1 | S_1 \leq \hat{S}_1(\mathbf{e}^{n+})] \right\}$$

and we let $\mathbf{e}^{0+} = \mathbf{e}$. In the expression above, $E_X[X_1 | X_1 \geq \hat{X}_1(\mathbf{e}^{n+})]$ is the expectation of the scrap value conditional on exiting the industry in state \mathbf{e}^{n+} and $E_S[S_1 | S_1 \leq \hat{S}_1(\mathbf{e}^{n+})]$ is the expectation of the setup cost conditional on entering. We refer the reader to the online Appendix for the counterpart $PS^{FB}(\mathbf{e})$ to $PS(\mathbf{e})$ under centralized exit and entry.

APPENDIX B. AGGRESSIVE AND ACCOMMODATIVE EQUILIBRIA

We note from the outset that any attempt to classify equilibria is fraught with difficulty because the different equilibria lie on a continuum and thus morph into each other in complicated ways as we vary the parameters of the model.

TABLE 3

	Unique equilibrium	Multiple equilibria		Multiple equilibria including accommodative equilibrium	
	49.09%	50.91%		21.98%	
		Best	Worst	Best	Worst
Aggressive	3.32%	43.36%	97.96%	5.84%	96.85%
Accommodative	83.40%	40.54%	1.36%	93.93%	3.15%
Unclassified	13.28%	16.10%	0.68%	0.22%	0.00%

Notes: Percentage of parameterizations with a unique equilibrium, multiple equilibria, and multiple equilibria that include an accommodative equilibrium and, within each of these, percentage of parameterizations at which the best or worst equilibrium is aggressive, accommodative, or unclassified. Unweighted.

Our definition of an aggressive equilibrium hones in on a trench in the pricing decision, and our definition of an accommodative equilibrium on a lack of exit from a duopolistic industry.

DEFINITION 1: An equilibrium is aggressive if $p_1(\mathbf{e}) < p_1(e_1, e_2 + 1)$, $p_2(\mathbf{e}) < p_2(e_1, e_2 + 1)$, and $\phi_2(\mathbf{e}) > \phi_2(e_1, e_2 + 1)$ for some state $\mathbf{e} > (0, 0)$ with $e_1 > e_2$.

DEFINITION 2: An equilibrium is accommodative if $\phi_1(\mathbf{e}) = \phi_2(\mathbf{e}) = 0$ for all states $\mathbf{e} > (0, 0)$.

These definitions are not exhaustive. The percentage of equilibria classified as aggressive is 96.88 percent, the percentage of equilibria classified as accommodative is 1.99 percent, and the percentage of unclassified equilibria is 1.13 percent. Our computations always led to a unique accommodative equilibrium but often to multiple aggressive equilibria at a given parameterization.

An aggressive equilibrium exists if the degree of product differentiation σ is sufficiently low. As competition in a duopolistic industry becomes fiercer, monopolizing the industry becomes more attractive. Other factors contributing to the existence of an aggressive equilibrium are a high expected scrap value \bar{X} and a high price of the outside good p_0 . The first factor makes it easier for a firm to induce its rival to exit the industry and the second makes monopolizing the industry more attractive by allowing the surviving firm to charge a higher price. Conversely, an accommodative equilibrium exists if the degree of product differentiation σ is sufficiently high. Other factors contributing to the existence of an accommodative equilibrium are a low expected scrap value \bar{X} and a low price of the outside good p_0 .

Our definitions of aggressive and accommodative equilibria are linked, however imperfectly, to worst, respectively, best equilibria. Table 3 shows that 49.09 percent of parameterizations have a unique equilibrium and that the equilibrium is classified as accommodative in 83.40 percent of these parameterizations. Moreover, 50.91 percent of parameterizations have multiple equilibria and the worst equilibrium is classified as aggressive in 97.96 percent of these parameterizations. However, the best equilibrium is classified as accommodative in 40.54 percent of these parameterizations and as aggressive in 43.36 percent of these parameterizations. This far from perfect link appears to reflect that our definition of an accommodative equilibrium is quite stringent. While there is an aggressive equilibrium at all parameterizations

with multiple equilibria, there may not be an accommodative equilibrium. The right-most column of Table 3 restricts attention to parameterizations with multiple equilibria that include an accommodative equilibrium. The best equilibrium is classified as accommodative in 93.93 percent of these parameterizations and the worst equilibrium is classified as aggressive in 96.85 percent of these parameterizations.

APPENDIX C. PROOFS

PROOF OF PROPOSITION 1:

Define the sum of consumer surplus and static profit to be

$$\Phi(p) = CS(p,p) + 2D_1(p,p)(p - c(e)).$$

Using this definition, in a symmetric state $\mathbf{e} = (e, e)$, where $e > 0$, we can write

$$CS^{FB}(\mathbf{e}) + \Pi^{FB}(\mathbf{e}) - (CS(\mathbf{e}) + \Pi(\mathbf{e})) = \Phi(p_1^{FB}(\mathbf{e})) - \Phi(p_1(\mathbf{e})).$$

We have

$$(5) \quad \Phi'(p) = -\frac{1}{\sigma}(p - c(e))D_0(p,p)(1 - D_0(p,p)),$$

$$(6) \quad \Phi''(p) = -\frac{1}{\sigma^2}((p - c(e))(1 - 2D_0(p,p)) + \sigma)D_0(p,p)(1 - D_0(p,p)).$$

Hence, $\Phi(p)$ is strictly quasi-concave in p and attains its maximum at $p = c(e)$. Thus, we obtain

$$(7) \quad CS^{FB}(\mathbf{e}) + \Pi^{FB}(\mathbf{e}) - (CS(\mathbf{e}) + \Pi(\mathbf{e})) \leq \Phi(c(e)) - \Phi(p_1(\mathbf{e})).$$

We bound the right-hand side of equation (7). Let \tilde{p} be such that $D_0(\tilde{p}, \tilde{p}) = 1/2$, so $1 - 2D_0(p,p) \geq 0$ for all $p \leq \tilde{p}$ because $D_0(p,p)$ increases in p . Equation (6) implies that $\Phi(p)$ is strictly concave in p over the interval $[c(e), \tilde{p}]$. This interval is non-empty: the assumption $p_0 \geq \kappa$ coupled with $\kappa \geq c(e)$ implies $D_0(c(e), c(e)) \leq D_0(\kappa, \kappa) \leq 1/3 < 1/2 = D_0(\tilde{p}, \tilde{p})$. As $D_0(p,p)$ increases in p , it must be that $c(e) < \tilde{p}$.

By assumption, $p_1(\mathbf{e}) \in [c(e), \tilde{p}]$. From Theorem 21.2 in Simon and Blume (1994) and equation (5) we therefore have

$$(8) \quad \Phi(c(e)) - \Phi(p_1(\mathbf{e})) \leq \Phi'(p_1(\mathbf{e}))(c(e) - p_1(\mathbf{e})) \\ = \frac{(p_1(\mathbf{e}) - c(e))^2}{\sigma} D_0(p_1(\mathbf{e}), p_1(\mathbf{e}))(1 - D_0(p_1(\mathbf{e}), p_1(\mathbf{e}))).$$

This establishes Proposition 1. ■

PROOF OF PROPOSITION 2:

Define the sum of consumer surplus and static profit to be

$$\Phi(p) = CS(p, \infty) + D_1(p, \infty)(p - c(e)).$$

Using this definition, in a state $\mathbf{e} = (e, 0)$, where $e > 0$, we can write

$$CS^{FB}(\mathbf{e}) + \Pi^{FB}(\mathbf{e}) - (CS(\mathbf{e}) + \Pi(\mathbf{e})) = \Phi(p_1^{FB}(\mathbf{e})) - \Phi(p_1(\mathbf{e})).$$

We have

$$\Phi'(p) = -\frac{1}{\sigma}(p - c(e))D_0(p, \infty)(1 - D_0(p, \infty)),$$

$$\Phi''(p) = -\frac{1}{\sigma^2}((p - c(e))(1 - 2D_0(p, \infty)) + \sigma)D_0(p, \infty)(1 - D_0(p, \infty)).$$

Hence, $\Phi(p)$ is strictly quasiconcave in p and attains its maximum at $p = c(e)$. Thus, we obtain

$$CS^{FB}(\mathbf{e}) + \Pi^{FB}(\mathbf{e}) - (CS(\mathbf{e}) + \Pi(\mathbf{e})) \leq \Phi(c(e)) - \Phi(p_1(\mathbf{e})),$$

where $\Phi(c(e)) = v - c(e) + \sigma \ln\left(1 + \exp\left(\frac{c(e) - p_0}{\sigma}\right)\right)$.

Note that, $p_1(\mathbf{e})$ is uniquely determined by the solution to the first-order condition (2); it can be written as

$$p_1(\mathbf{e}) = c(e) + \sigma(-y + 1 + W(\exp(x + y - 1))),$$

where $W(\cdot)$ is the Lambert W function¹⁵ and we define $x = \frac{p_0 - c(e)}{\sigma}$ and $y = (U_1(\mathbf{e}^{1+}) - U_1(\mathbf{e}))/\sigma$. Hence, using the identities $W(z) = z/\exp(W(z))$ for all z and $\ln W(z) = \ln z - W(z)$ for all $z > 0$,

$$\Phi(p_1(\mathbf{e})) = v - c(e) + \sigma\left(\ln\left(\frac{1 + W(\exp(x + y - 1))}{W(\exp(x + y - 1))}\right) + \frac{y}{1 + W(\exp(x + y - 1))} - 1\right)$$

and

$$\Phi(c(e)) = v - c(e) + \sigma \ln(1 + \exp(-x)).$$

Using the identity $\ln W(z) = \ln z - W(z)$ for all $z > 0$, it follows that

$$\begin{aligned} (9) \quad & \Phi(c(e)) - \Phi(p_1(\mathbf{e})) \\ &= \sigma\left(\ln(1 + \exp(-x)) - \ln\left(\frac{1 + W(\exp(x + y - 1))}{W(\exp(x + y - 1))}\right) - \frac{y}{1 + W(\exp(x + y - 1))} + 1\right) \end{aligned}$$

¹⁵The Lambert W function is defined by the equation $z = W(z)\exp(W(z))$ for all $z \in \mathbb{C}$.

$$(10) = \sigma \left(\ln(1 + \exp(x)) - \ln(1 + W(\exp(x + y - 1))) - W(\exp(x + y - 1)) + \frac{yW(\exp(x + y - 1))}{1 + W(\exp(x + y - 1))} \right).$$

In the remainder of the proof, we use that $W(z)$ is increasing in z , $W(0) = 0$, and the identity $W(z \exp(z)) = z$ for all $z \geq 0$.

Case 1: $y < 1 + \exp(x)$. We first show that if $y < 1 + \exp(x)$, then

$$(11) \quad \ln(1 + \exp(-x)) - \ln\left(\frac{1 + W(\exp(x + y - 1))}{W(\exp(x + y - 1))}\right) + 1 < 1.$$

To see this, note that equation (11) is equivalent to

$$\ln(1 + \exp(-x)) < \ln\left(\frac{1 + W(\exp(x + y - 1))}{W(\exp(x + y - 1))}\right) \Leftrightarrow \exp(x) > W(\exp(x + y - 1)).$$

If $y = 1 + \exp(x)$, then the right-hand side is $W(\exp(x + \exp(x))) = \exp(x)$. Moreover, because $W(z)$ is increasing in z , $\exp(x) > W(\exp(x + y - 1))$ for all $y < 1 + \exp(x)$.

Consider equation (9). From equation (11) it follows that

$$\Phi(c(e)) - \Phi(p_1(e)) < \sigma \left(1 - y \frac{1}{1 + W(\exp(x + y - 1))} \right).$$

Moreover, $0 < \frac{1}{1 + W(\exp(x + y - 1))} < 1$. Therefore, if $y < 0$, then

$$(12) \quad \Phi(c(e)) - \Phi(p_1(e)) < \sigma(1 + |y|)$$

and, if $y \geq 0$, then

$$(13) \quad \Phi(c(e)) - \Phi(p_1(e)) < \sigma.$$

Case 2: $y \geq 1 + \exp(x)$. We first show that if $y \geq 1 + \exp(x)$, then

$$(14) \quad \ln(1 + \exp(x)) - \ln(1 + W(\exp(x + y - 1))) - W(\exp(x + y - 1)) < 1.$$

To see this, note that

$$\begin{aligned} & \ln(1 + \exp(x)) - \ln(1 + W(\exp(x + y - 1))) - W(\exp(x + y - 1)) \\ & \leq \ln(1 + \exp(x)) - \ln(1 + W(\exp(x + \exp(x)))) - W(\exp(x + \exp(x))) \\ & = \ln(1 + \exp(x)) - \ln(1 + \exp(x)) - \exp(x) = -\exp(x) < 1. \end{aligned}$$

Consider equation (10). From equation (14) it follows that

$$\Phi(c(e)) - \Phi(p_1(e)) < \sigma \left(1 + y \frac{W(\exp(x + y - 1))}{1 + W(\exp(x + y - 1))} \right).$$

Moreover, $0 < \frac{W(\exp(x + y - 1))}{1 + W(\exp(x + y - 1))} < 1$. Because $y \geq 1 + \exp(x) > 0$, we have

$$(15) \quad \Phi(c(e)) - \Phi(p_1(e)) < \sigma(1 + y).$$

Collecting equations (12), (13), and (15) establishes Proposition 2. ■

REFERENCES

- Athey, Susan, and Armin Schmutzler. 2001. "Investment and Market Dominance." *RAND Journal of Economics* 32 (1): 1–26.
- Bergemann, Dirk, and Juuso Välimäki. 1996. "Learning and Strategic Pricing." *Econometrica* 64 (5): 1125–49.
- Bergemann, Dirk, and Juuso Välimäki. 1997. "Market Diffusion with Two-Sided Learning." *RAND Journal of Economics* 28 (4): 773–95.
- Bergemann, Dirk, and Juuso Välimäki. 2006. "Dynamic Price Competition." *Journal of Economic Theory* 127 (1): 232–63.
- Besanko, David, Ulrich Doraszelski, and Yaroslav Kryukov. 2014. "The Economics of Predation: What Drives Pricing When There Is Learning-by-Doing?" *American Economic Review* 104 (3): 868–97.
- Besanko, David, Ulrich Doraszelski, and Yaroslav Kryukov. 2019. "How Efficient Is Dynamic Competition? The Case of Price as Investment: Dataset." *American Economic Review*. <https://doi.org/10.1257/aer.20180131>.
- Besanko, David, Ulrich Doraszelski, Yaroslav Kryukov, and Mark Satterthwaite. 2010. "Learning-by-Doing, Organizational Forgetting, and Industry Dynamics." *Econometrica* 78 (2): 453–508.
- Biglaiser, Gary, and Jacques Crémer. 2016. "The Value of Incumbency for Heterogeneous Platforms." Toulouse School of Economics Working Paper 16-630.
- Bolton, Patrick, and Joseph Farrell. 1990. "Decentralization, Duplication, and Delay." *Journal of Political Economy* 98 (4): 803–26.
- Bolton, Patrick, and Christopher Harris. 1999. "Strategic Experimentation." *Econometrica* 67 (2): 349–74.
- Bulow, Jeremy, and Paul Klemperer. 1999. "The Generalized War of Attrition." *American Economic Review* 89 (1): 175–89.
- Cabral, Luís. 2011. "Dynamic Price Competition with Network Effects." *Review of Economic Studies* 78 (1): 83–111.
- Cabral, Luís M. B., and Michael H. Riordan. 1994. "The Learning Curve, Market Dominance, and Predatory Pricing." *Econometrica* 62 (5): 1115–40.
- Chen, Jiawei. 2011. "How Do Switching Costs Affect Market Concentration and Prices in Network Industries?" Unpublished.
- Chen, Jiawei, Ulrich Doraszelski, and Joseph E. Harrington. 2009. "Avoiding Market Dominance: Product Compatibility in Markets with Network Effects." *RAND Journal of Economics* 40 (3): 455–85.
- Ching, Andrew T. 2010. "A Dynamic Oligopoly Structural Model for the Prescription Drug Market after Patient Expiration." *International Economic Review* 51 (4): 1175–1207.
- Dasgupta, Partha, and Joseph Stiglitz. 1988. "Learning-by-Doing, Market Structure and Industrial and Trade Policies." *Oxford Economic Papers* 40 (2): 246–68.
- Doraszelski, Ulrich, and Mark Satterthwaite. 2010. "Computable Markov-Perfect Industry Dynamics." *RAND Journal of Economics* 41 (2): 215–43.
- Dubé, Jean-Pierre, Günter J. Hitsch, and Pradeep K. Chintagunta. 2010. "Tipping and Concentration in Markets with Indirect Network Effects." *Marketing Science* 29 (2): 216–49.
- Dubé, Jean-Pierre, Günter J. Hitsch, and Peter E. Rossi. 2009. "Do Switching Costs Make Markets Less Competitive?" *Journal of Marketing Research* 46 (4): 435–45.

- Halaburda, Hanna, Bruno Jullien, and Yaron Yehezkel.** 2016. "Dynamic Competition with Network Externalities: Why History Matters." Toulouse School of Economics Working Paper TSE-636.
- Harberger, Arnold C.** 1954. "Monopoly and Resource Allocation." *American Economic Review* 44 (2): 77–87.
- Mitchell, Matthew F., and Andrzej Skrzypacz.** 2006. "Network Externalities and Long-Run Market Shares." *Economic Theory* 29 (3): 621–48.
- Posner, Richard A.** 1975. "The Social Costs of Monopoly and Regulation." *Journal of Political Economy* 83 (4): 807–27.
- Rob, Rafael.** 1991. "Learning and Capacity Expansion under Demand Uncertainty." *Review of Economic Studies* 58 (4): 655–75.
- Simon, Carl P., and Lawrence Blume.** 1994. *Mathematics for Economists*. New York: W.W. Norton & Company.
- Smith, J. Maynard.** 1974. "Theory of Games and the Evolution of Animal Conflicts." *Journal of Theoretical Biology* 47 (1): 209–21.
- Tirole, Jean.** 1988. *The Theory of Industrial Organization*. Cambridge, MA: MIT Press.