RAND Journal of Economics
Vol. 41, No. 2, Summer 2010
pp. 215–243

Computable Markov-perfect industry dynamics

Ulrich Doraszelski*

and

Mark Satterthwaite**

We provide a general model of dynamic competition in an oligopolistic industry with investment, entry, and exit. To ensure that there exists a computationally tractable Markov-perfect equilibrium, we introduce firm heterogeneity in the form of randomly drawn, privately known scrap values and setup costs into the model. Our game of incomplete information always has an equilibrium in cutoff entry/exit strategies. In contrast, the existence of an equilibrium in the Ericson and Pakes' model of industry dynamics requires admissibility of mixed entry/exit strategies, contrary to the assertion in their article, that existing algorithms cannot cope with. In addition, we provide a condition on the model's primitives that ensures that the equilibrium is in pure investment strategies. Building on this basic existence result, we first show that a symmetric equilibrium exists under appropriate assumptions on the model's primitives. Second, we show that, as the distribution of the random scrap values/setup costs becomes degenerate, equilibria in cutoff entry/exit strategies converge to equilibria in mixed entry/exit strategies of the game of complete information.

1. Introduction

■ Building on the seminal work of Maskin and Tirole (1987, 1988a, 1988b), the industrial organization literature has made considerable progress over the past few years in analyzing industry dynamics. In an important article, Ericson and Pakes (1995) provide a computable model of dynamic competition in an oligopolistic industry with investment, entry, and exit. Their framework is a valuable addition to economists' toolkits. Its applications to date have yielded numerous novel insights and it provides a starting point for ongoing research in industrial

*Harvard University and CEPR; doraszelski@harvard.edu.
**Northwestern University; m-satterthwaite@kellogg.northwestern.edu.
We are indebted to Lanier Benkard, David Besanko, Michaela Draganska, Dino Gerardi, Gautam Gowrisankaran, Ken Judd, Patricia Langohr, Andy Skrzypacz, and Eilon Solan for useful discussions. The comments of the editor, Joe Harrington, and two anonymous referees greatly helped to improve the article. Satterthwaite acknowledges gratefully that this material is based upon work supported by the National Science Foundation under grant no. 0121541 and by the General Motors Research Center for Strategy in Management at Kellogg. Doraszelski is grateful for the hospitality of the Hoover Institution during the academic year 2006/07 and financial support from the National Science Foundation under grant no. 0615615.

organization and other fields (see Doraszelski and Pakes, 2007 for a survey). More recently, Aguirregabiria and Mira (2007), Bajari, Benkard, and Levin (2007), Pakes, Ostrovsky, and Berry (2007), and Pesendorfer and Schmidt-Dengler (2008) have developed estimation procedures that allow the researcher to recover the primitives that underlie the dynamic industry equilibrium. Consequently, it is now possible to take these models to the data with the goal of conducting counterfactual experiments and policy analyses (e.g., Gowrisankaran and Town, 1997; Jofre-Bonet and Pesendorfer, 2003; Benkard, 2004; Beresteanu and Ellickson, 2006; Collard-Wexler, 2006; Ryan, 2006).

To achieve this goal, the researcher has to be able to compute the stationary Markov-perfect equilibrium using the estimated primitives. This, in turn, requires that an equilibrium exists. Unfortunately, existence cannot be guaranteed under the conditions in Ericson and Pakes (1995). Moreover, existence by itself is not enough for two reasons. First, contrary to the assertion in their article, the existence of an equilibrium in the Ericson and Pakes' (1995) model of industry dynamics requires admissibility of mixed strategies over discrete actions such as entry and exit. But computing mixed strategies poses a formidable challenge (even in the context of finite games; see McKelvey and McLennan, 1996 for a survey). Second, the state space of the model explodes in the number of firms and quickly overwhelms current computational capabilities. An important means of mitigating this "curse of dimensionality" is to impose symmetry restrictions. For these reasons, computational tractability requires existence of a symmetric equilibrium in pure strategies.

Our goal in this article is to modify the Ericson and Pakes' (1995) model just enough to ensure that there exists for it a stationary Markov-perfect equilibrium that is computable in both theory and practice.[1] In doing so, we have to resolve three difficulties that we now discuss in detail.

□ **Cutoff entry/exit strategies.** In the Ericson and Pakes' (1995) model, incumbent firms decide in each period whether to remain in the industry and potential entrants decide whether to enter the industry. But the existence of an equilibrium cannot be ensured without allowing firms to randomize, in one way or another, over these discrete actions. Because Ericson and Pakes (1995) do not provide for such mixing, a simple example suffices to show that their claim of existence cannot possibly be correct (see Section 3).[2]

To eliminate the need for mixed entry/exit strategies without jeopardizing existence, we extend Harsanyi's (1973) technique for purifying mixed-strategy Nash equilibria of static games to Markov-perfect equilibria of dynamic stochastic games, and assume that at the beginning of each period each potential entrant is assigned a random setup cost payable upon entry and each incumbent firm is assigned a random scrap value received upon exit. Setup costs/scrap values are privately known, that is, whereas a firm learns its own setup cost/scrap value prior to making its decisions, its rivals' setup costs/scrap values remain unknown to it. Adding firm heterogeneity in the form of these randomly drawn, privately known setup costs/scrap values leads to a game of incomplete information. This game always has an equilibrium in cutoff entry/exit strategies that existing algorithms—notably Pakes and McGuire (1994, 2001)—can handle after minor changes. Although a firm formally follows a pure strategy in making its entry/exit decision, the dependence of its entry/exit decision on its randomly drawn, privately known setup cost/scrap value implies that its rivals perceive the firm *as though* it were following a mixed strategy. Note that random

[1] Given that an equilibrium exists, an important question is whether or not it is unique. In the online appendix to this article, we show that multiplicity may be an issue in Ericson and Pakes' (1995) framework even if symmetry restrictions are imposed by providing three examples of multiple symmetric equilibria.

[2] The game-theoretic literature has, of course, recognized the importance of randomization, but relies on computationally intractable mixed strategies (see Mertens, 2002 for a survey). Strictly speaking, the existence theorems in the extant literature are not even applicable because they cover dynamic stochastic games with either discrete (e.g., Fink, 1964; Sobel, 1971; Maskin and Tirole, 2001) or continuous actions (e.g., Federgruen, 1978; Whitt, 1980), whereas Ericson and Pakes' (1995) model combines discrete entry/exit decisions with continuous investment decisions.

setup costs/scrap values can substitute for mixed entry/exit strategies only if they are privately known. If they were publicly observed, then its rivals could infer with certainty whether or not the firm will enter/exit the industry. In this manner, Harsanyi's (1973) insight that a perturbed game of incomplete information can purify the mixed-strategy equilibria of an underlying game of complete information enables us to settle the first and perhaps central difficulty in devising a computationally tractable model.[3]

Over the years, the idea of using random setup costs/scrap values instead of mixed entry/exit strategies has become part of the folklore in the literature following Ericson and Pakes (1995). Pakes and McGuire (1994) suggest treating a potential entrant's setup cost as a random variable to overcome convergence problems in their algorithm. Gowrisankaran (1999a) has an informal but very clear discussion of how randomization can resolve existence issues whenever entry, exit, or mergers are allowed. Nevertheless, neither that article nor Gowrisankaran (1995) provide a precise, rigorous, and reasonably general statement of how randomization can be inserted into the Ericson and Pakes' (1995) model so as to guarantee existence.

More recently, in independent work, Aguirregabiria and Mira (2007) and Pesendorfer and Schmidt-Dengler (2008) use randomly drawn, privately known shocks to establish the existence of a Markov-perfect equilibrium in general dynamic stochastic games with finite state and action spaces.[4] The primary difference between their discrete-choice frameworks and our model is that we allow for continuous as well as discrete actions. Because discretizing continuous actions tends to complicate both the estimation and computation of the model, most applied work treats actions such as advertising (e.g., Doraszelski and Markovich, 2007), investment (e.g., Besanko and Doraszelski, 2004; Beresteanu and Ellickson, 2006; Ryan, 2006), and price (e.g., Besanko, Doraszelski, Kryukov, and Satterthwaite, 2010) as continuous variables. The arguments we develop here can be used to guarantee existence in all these cases.

The existing literature also leaves open the important question of whether the "trick" of using random setup costs/scrap values changes the nature of strategic interactions among firms. We show that, as the distribution of the random scrap values/setup costs becomes degenerate, an equilibrium in cutoff entry/exit strategies of the incomplete-information game converges to an equilibrium in mixed entry/exit strategies of the complete-information game (see Section 7). Hence, the addition of random scrap values/setup costs does not change the nature of strategic interactions among firms. An immediate consequence of our convergence result is that there exists an equilibrium in the Ericson and Pakes' (1995) model, provided that mixed entry/exit strategies are admissible.

□ **Pure investment strategies.** In addition to deciding whether to remain in the industry, incumbent firms also decide how much to invest in each period in the Ericson and Pakes' (1995) model. Because mixed strategies over continuous actions are impractical to compute, the second difficulty is to ensure pure investment strategies. One way to forestall the possibility of mixing is to make sure that a firm's optimal investment level is always unique. To achieve this, we define a class of transition functions—functions which specify how firms' investment decisions affect the industry's state-to-state transitions—that we call unique investment choice (UIC) admissible and prove that if the transition function is UIC admissible, then a firm's investment choice is indeed uniquely determined (see Section 5). UIC admissibility is an easily verifiable condition on the model's primitives and is not overly limiting. Indeed, although the transition functions used in the vast majority of applications of Ericson and Pakes' (1995) framework are UIC admissible, they all restrict a firm to transit to immediately adjacent states. Our condition establishes that this is unnecessary, and we show how to specify more general UIC admissible transition functions.

[3] There is an interesting parallel between our article, which puts "noise" in the payoffs, and papers that put "noise" in the state-to-state transitions in order to overcome existence problems in dynamic stochastic games with continuous state spaces; see the excellent summary of this literature in Chakrabarti (1999).

[4] The working paper versions of all three articles were initially circulated between May 2003 and September 2004; see Pesendorfer and Schmidt-Dengler (2003), Doraszelski and Satterthwaite (2003), and Aguirregabiria and Mira (2004).

In subsequent work, Escobar (2008) establishes the existence of a Markov-perfect equilibrium in pure strategies in a general dynamic stochastic game with a countable state space and a continuum of actions. He follows an approach similar to ours by first proving existence under the assumption that a player's best reply is convex for any value of continued play and then characterizing the class of per-period payoffs and transition functions that ensure that this is indeed the case. Because a unique best reply is a special case of a convex best reply, his condition is more general than ours and may be applied to games with continuous actions other than the investment decisions in the Ericson and Pakes' (1995) model.[5]

□ **Symmetry.** The third and final difficulty in devising a computationally tractable model is to ensure that the equilibrium is not only in pure strategies but is also symmetric. We show that this is the case under appropriate assumptions on the model's primitives (see Section 6). Symmetry is important because it eases the computational burden considerably. Instead of having to compute value and policy functions for *all* firms, under symmetry it suffices to compute value and policy functions for *one* firm. In addition, symmetry reduces the size of the state space on which these functions are defined. Besides its computational advantages, a symmetric equilibrium is an especially convincing solution concept in models of dynamic competition with entry and exit because there is often no reason why a particular entrant should be different from any other entrant. Rather, firm heterogeneity arises endogenously from the idiosyncratic outcomes that the *ex ante* identical firms realize from their investments.

Resolving these difficulties allows us to fulfill our goal of establishing that there always exists a stationary Markov-perfect equilibrium that is symmetric and in pure strategies. A further goal of this article is to provide a step-by-step guide to formulating models of dynamic industry equilibrium that is detailed enough to allow the reader to easily adapt its techniques to models that are tailored to specific industries. We hope that such a guide enables others to construct their models with the confidence that if their algorithm fails to converge, it is a computational problem, not a poorly specified model for which no equilibrium exists.

The plan of the article is as follows. We develop the model in Section 2. In Section 3, we provide simple examples to illustrate the key themes of the subsequent analysis. We turn to the analysis of the full model in Sections 4–7. Section 8 concludes.

2. Model

■ We study the evolution of an industry with heterogeneous firms. The model is dynamic, time is discrete, and the horizon is infinite. There are two groups of firms, incumbent firms and potential entrants. An incumbent firm has to decide each period whether to remain in the industry and, if so, how much to invest. A potential entrant has to decide whether to enter the industry and, if so, how much to invest. Once these decisions are made, product market competition takes place.

Our model accounts for firm heterogeneity in two ways. First, we encode all characteristics that are relevant to a firm's profit from product market competition (e.g., production capacity, cost structure, or product quality) in its "state." A firm is able to change its state over time through investment. Although a higher investment today is no guarantee for a more favorable state tomorrow, it does ensure a more favorable distribution over future states. By acknowledging that a firm's transition from one state to another is subject to an idiosyncratic shock, our model allows for variability in the fortunes of firms even if they carry out identical strategies. Second,

[5] Other work on existence in pure strategies includes Dutta and Sundaram (1992) (resource extraction games), Amir (1996) (capital accumulation games), Curtat (1996) and Nowak (2007) (supermodular games), and Horst (2005) (games with weak interactions among players). Chakrabarti (2003) studies games with a continuum of players in which the per-period payoffs and the transition density function depend only on the *average* response of the players.

to account for differences in opportunity costs across firms, we assume that incumbents have random scrap values (received upon exit) and that entrants have random setup costs (payable upon entry). Because a firm's particular circumstances change over time, we model scrap values and setup costs as being drawn anew each period.

□ **States and firms.** Let N denote the number of firms. Firm n is described by its state $\omega_n \in \Omega$, where $\Omega = \{1, \dots, M, M+1\}$ is its set of possible states. States $1, \dots, M$ describe an active firm whereas state $M+1$ identifies the firm as being inactive.[6] At any point in time the industry is completely characterized by the list of firms' states $\omega = (\omega_1, \dots, \omega_N) \in S$, where $S = \Omega^N$ is the state space.[7] We refer to ω_n as the state of firm n and to ω as the state of the industry.

If N^* is the number of incumbent firms (i.e., active firms), then there are $N - N^*$ potential entrants (i.e., inactive firms). Thus, once an incumbent firm exits the industry, a potential entrant automatically takes its "slot" and has to decide whether or not to enter the industry.[8] Potential entrants are drawn from a large pool. They are short lived and base their entry decisions on the net present value of entering today; potential entrants do not take the option value of delaying entry into account. In contrast, incumbent firms are long lived and solve intertemporal maximization problems to reach their exit decisions. They discount future payoffs using a discount factor of β.

□ **Timing.** In each period the sequence of events is as follows:

(i) Incumbent firms learn their scrap value and decide on exit and investment. Potential entrants learn their setup cost and decide on entry and investment.
(ii) Incumbent firms compete in the product market.
(iii) Exit and entry decisions are implemented.
(iv) The investment decisions of the remaining incumbents and new entrants are carried out and their uncertain outcomes are realized.

Throughout, we use ω to denote the state of the industry at the beginning of the period and ω' to denote its state at the end of the period after the state-to-state transitions are realized. Firms observe the state at the beginning of the period as well as the outcomes of the entry, exit, and investment decisions during the period.

Whereas the entry, exit, and investment decisions are made simultaneously, we assume that an incumbent's investment decision is carried out only if it remains in the industry. Similarly, we assume that an entrant's investment decision is carried out only if it enters the industry. It follows that an optimizing incumbent firm will choose its investment at the beginning of each period under the presumption that it does not exit this period, and an optimizing potential entrant will do so under the presumption that it enters the industry.

□ **Incumbent firms.** Suppose $\omega_n \neq M+1$ and consider incumbent firm n. We assume that at the beginning of each period each incumbent firm draws a random scrap value ϕ_n from a distribution $F(\cdot)$ with expectation $E(\phi_n) = \phi$.[9] Scrap values are independently and identically distributed across firms and periods. Incumbent firm n learns its scrap value ϕ_n prior to making its exit and investment decisions, but the scrap values of its rivals remain unknown to it. Let

[6] This formulation allows firms to differ from each other in more than one dimension. Suppose that a firm is characterized by its capacity and its marginal cost of production. If there are M_1 levels of capacity and M_2 levels of cost, then each of the $M = M_1 M_2$ possible combinations of capacity and cost defines a state.

[7] Time-varying characteristics of the competitive environment are easily added to the description of the industry. Besanko and Doraszelski (2004), for example, add a demand state to the list of firms' states in order to study the effects of demand growth and demand cycles on capacity dynamics.

[8] Limiting the number of potential entrants to $N - N^*$ is not innocuous. Increasing $N - N^*$ by increasing N exacerbates the coordination problem that potential entrants face.

[9] It is straightforward to allow firm n's scrap value ϕ_n to vary systematically with its state ω_n by replacing $F(\cdot)$ by $F_{\omega_n}(\cdot)$.

$\chi_n(\omega, \phi_n) = 1$ indicate that the decision of incumbent firm n, who has drawn scrap value ϕ_n, is to remain in the industry in state ω and let $\chi_n(\omega, \phi_n) = 0$ indicate that its decision is to exit the industry, collect the scrap value ϕ_n, and perish. Because this decision is conditioned on its private ϕ_n, it is a random variable from the perspective of other firms. We use $\xi_n(\omega) = \int \chi_n(\omega, \phi_n) dF(\phi_n)$ to denote the probability that incumbent firm n remains in the industry in state ω.

This is the first place where our model diverges from Ericson and Pakes' (1995), who assume that scrap values are constant across firms and periods. As we show in Section 3, deterministic scrap values raise serious existence issues. In the limit, however, as the distribution of ϕ_n becomes degenerate, our model collapses to theirs.

If an incumbent remains in the industry, it competes in the product market. Let $\pi_n(\omega)$ denote the current profit of incumbent firm n from product market competition in state ω. We stipulate that $\pi_n(\cdot)$ is a reduced-form profit function that fully incorporates the nature of product market competition in the industry. In addition to receiving a profit, the incumbent incurs the investment $x_n(\omega) \in [0, \bar{x}]$ that it decided on at the beginning of the period and moves from state ω_n to state $\omega_n' \neq M + 1$ in accordance with the transition probabilities specified below.

□ **Potential entrants.** Suppose that $\omega_n = M + 1$ and consider potential entrant n. We assume that at the beginning of each period each potential entrant draws a random setup cost ϕ_n^e from a distribution $F^e(\cdot)$ with expectation $\mathrm{E}(\phi_n^e) = \phi^e$. Like scrap values, setup costs are independently and identically distributed across firms and periods, and are private to each firm. If potential entrant n enters the industry, it incurs the setup cost ϕ_n^e. If it stays out, it receives nothing and perishes. We use $\chi_n^e(\omega, \phi_n^e) = 1$ to indicate that the decision of potential entrant n, who has drawn setup cost ϕ_n^e, is to enter the industry in state ω and $\chi_n^e(\omega, \phi_n^e) = 0$ to indicate that its decision is to stay out. From the point of view of other firms, $\xi_n^e(\omega) = \int \chi_n^e(\omega, \phi_n^e) dF^e(\phi_n^e)$ denotes the probability that potential entrant n enters the industry in state ω.

Unlike an incumbent, the entrant does not compete in the product market. Instead, it undergoes a setup period upon committing to entry. The entrant incurs its previously chosen investment $x_n^e(\omega) \in [0, \bar{x}^e]$ and moves to state $\omega_n' \neq M + 1$. Hence, at the end of the setup period, the entrant becomes an incumbent.

This is the second place where we generalize the Ericson and Pakes' (1995) model. They assume that, unlike exit decisions, entry decisions are made sequentially. We assume that entry decisions are made simultaneously, thus allowing more than one firm per period to enter the industry in an uncoordinated fashion. We also allow the potential entrant to make an initial investment in order to improve the odds that it enters the industry in a more favorable state. This contrasts with Ericson and Pakes (1995), where the entrant is randomly assigned to an arbitrary position and thus has no control over its initial position within the industry.[10]

We make these two changes because industry evolution frequently takes the form of a preemption race (e.g., Fudenberg, Gilbert, Stiglitz, and Tirole, 1983; Harris and Vickers, 1987; Besanko and Doraszelski, 2004; Doraszelski and Markovich, 2007). During such a race, firms invest heavily as long as they are neck and neck. But once one of the firms manages to pull ahead, the lagging firms "give up," thereby allowing the leading firm to attain a dominant position. In a preemption race, an early entrant has a head start over a late entrant, so an imposed order of entry may prove to be decisive for the structure of the industry. Moreover, denying an entrant control over its initial position within the industry makes it all the harder to "catch up." Our specification of the entry process does not suffer from these drawbacks and makes the model more realistic by endogenizing the intensity of entry activity. As an additional benefit, our parallel treatment of entry and exit as well as incumbents' and entrants' investment decisions simplifies the model's exposition and eases the computational burden.

[10] We may nest their formulation by setting $\bar{x}^e = 0$.

□ **Notation.** In what follows, we identify the nth incumbent firm with firm n in states $\omega_n \neq M+1$ and the nth potential entrant with firm n in state $\omega_n = M+1$. That is, we define

$$\chi_n^e(\omega_1,\ldots,\omega_{n-1},\omega_n,\omega_{n+1},\ldots,\omega_N,\phi^e) = \chi_n(\omega_1,\ldots,\omega_{n-1},M+1,\omega_{n+1},\ldots,\omega_N,\phi^e),$$
$$\xi_n^e(\omega_1,\ldots,\omega_{n-1},\omega_n,\omega_{n+1},\ldots,\omega_N) = \xi_n(\omega_1,\ldots,\omega_{n-1},M+1,\omega_{n+1},\ldots,\omega_N),$$
$$x_n^e(\omega_1,\ldots,\omega_{n-1},\omega_n,\omega_{n+1},\ldots,\omega_N) = x_n(\omega_1,\ldots,\omega_{n-1},M+1,\omega_{n+1},\ldots,\omega_N).$$

Because ω_n indicates whether firm n is an incumbent firm or a potential entrant, we henceforth omit the superscript e to distinguish entrants from incumbents.

□ **Transition probabilities.** The probability that the industry transits from today's state ω to tomorrow's state ω' is determined jointly by the investment decisions of the incumbent firms that remain in the industry and the potential entrants that enter the industry. Formally, the transition probabilities are encoded in the transition function $P : S^2 \times \{0,1\}^N \times [0, \max\{\bar{x}, \bar{x}^e\}]^N \rightarrow [0,1]$. Thus, $P(\omega', \omega, \chi(\omega,\phi), x(\omega))$ is the probability that the industry moves from state ω to state ω' given that firms' exit and entry decisions are $\chi(\omega,\phi) = (\chi_1(\omega,\phi_1),\ldots,\chi_N(\omega,\phi_N))$ and their investment decisions are $x(\omega) = (x_1(\omega),\ldots,x_N(\omega))$.[11] Necessarily, $P(\omega', \omega, \chi(\omega,\phi), x(\omega)) \geq 0$ and $\sum_{\omega' \in S} P(\omega', \omega, \chi(\omega,\phi), x(\omega)) = 1$.

In the special case of independent transitions, the transition function $P(\cdot)$ can be factored as

$$\prod_{n=1,\ldots,N} P_n(\omega_n', \omega_n, \chi_n(\omega,\phi_n), x_n(\omega)),$$

where $P_n(\cdot)$ gives the probability that firm n transits from state ω_n to state ω_n' conditional on its exit or entry decision being $\chi_n(\omega,\phi_n)$ and its investment decision being $x_n(\omega)$. In general, however, transitions need not be independent across firms. Independence is violated, for example, in the presence of demand or cost shocks that are common to firms or in the presence of externalities.

Because a firm's scrap value or setup cost is private information, its exit or entry decision is a random variable from the perspective of an outside observer. The outside observer thus has to "integrate out" over all possible realizations of firms' exit and entry decisions to obtain the probability that the industry transits from state ω to state ω':

$$\int \cdots \int P(\omega', \omega, \chi(\omega,\phi), x(\omega)) \prod_{\substack{n=1,\ldots,N,\\ \omega_n \neq M+1}} dF(\phi_n) \prod_{\substack{n=1,\ldots,N,\\ \omega_n = M+1}} dF^e(\phi_n^e)$$
$$= \sum_{\iota \in \{0,1\}^N} \left[P(\omega', \omega, \iota, x(\omega)) \prod_{n=1,\ldots,N} \xi_n(\omega)^{\iota_n} (1-\xi_n(\omega))^{1-\iota_n} \right]. \quad (1)$$

To see this, recall that scrap values and setup costs are independently distributed across firms. Because, from the point of view of other firms, the probability that incumbent firm n remains in the industry in state ω is $\xi_n(\omega) = \int \chi_n(\omega,\phi_n) dF(\phi_n)$ and the probability that potential entrant n enters the industry is $\xi_n(\omega) = \int \chi_n(\omega,\phi_n^e) dF^e(\phi_n^e)$, a particular realization $\iota = (\iota_1,\ldots,\iota_N)$ of firms' exit and entry decisions occurs with probability $\prod_{n=1,\ldots,N} \xi_n(\omega)^{\iota_n}(1-\xi_n(\omega))^{1-\iota_n}$. In this manner, equation (1) results from conditioning on all possible realizations of firms' exit and entry decisions ι.

The crucial implication of equation (1) is that the probability of a transition from state ω to state ω' hinges on the exit and entry probabilities $\xi(\omega)$. Thus, when forming an expectation over the industry's future state, a firm does not need to know the entire exit and entry rules $\chi_{-n}(\omega,\cdot)$ of its rivals; rather, it suffices to know their exit and entry probabilities $\xi_{-n}(\omega)$.

[11] Given our notational convention, if $\omega_n = M+1$ so that firm n is a potential entrant, then we interpret $\chi_n(\omega,\phi_n)$ as $\chi_n^e(\omega,\phi_n^e)$, the decision of potential entrant n, who has drawn setup cost ϕ_n^e, to enter the industry in state ω, and we similarly interpret $x_n(\omega)$ as $x_n^e(\omega)$.

□ **An incumbent's problem.** Suppose that the industry is in state ω with $\omega_n \neq M+1$. Incumbent firm n solves an intertemporal maximization problem to reach its exit and investment decisions. Let $V_n(\omega, \phi_n)$ denote the expected net present value of all future cash flows to incumbent firm n, computed under the presumption that firms behave optimally, when the industry is in state ω and the incumbent has drawn scrap value ϕ_n. Note that its scrap value is part of the payoff-relevant characteristics of the incumbent firm. This is rather obvious: an incumbent that can sell off its assets for one dollar may behave very differently from an otherwise identical incumbent that can sell off its assets for one million dollars. Hence, once incumbent firm n has learned its scrap value ϕ_n, its decisions and thus also the expected net present value of its future cash flows, $V_n(\omega, \phi_n)$, depend on it. Unlike deterministic scrap values, random scrap values are part of the state space of the game. This is undesirable from a computational perspective because the computational burden is increasing with the size of the state space. Fortunately, as we show below, integrating out over the random scrap values eliminates their disadvantage but preserves their advantage for ensuring the existence of an equilibrium.

$V_n(\omega, \phi_n)$ is defined recursively by the solution to the following Bellman equation,

$$V_n(\omega, \phi_n) = \sup_{\substack{\tilde{\chi}_n(\omega,\phi_n) \in \{0,1\},\\ \tilde{x}_n(\omega) \in [0,\bar{x}]}} \pi_n(\omega) + (1 - \tilde{\chi}_n(\omega, \phi_n))\phi_n + \tilde{\chi}_n(\omega, \phi_n) \times \{-\tilde{x}_n(\omega) + \beta \mathrm{E}\{V_n(\omega')|\omega, \omega'_n \neq M+1, \tilde{x}_n(\omega), \xi_{-n}(\omega), x_{-n}(\omega)\}\}, \qquad (2)$$

where, with an overloading of notation, $V_n(\omega) = \int V_n(\omega, \phi_n)dF(\phi_n)$ is the expected value function. Note that whereas $V_n(\omega, \phi_n)$ is the value function *after* the firm has drawn its scrap value, $V_n(\omega)$ is the expected value function, that is, the value function *before* the firm has drawn its scrap value. The right-hand side of the Bellman equation is composed of the incumbent's profit from product market competition $\pi_n(\omega)$ and, depending on the exit decision $\tilde{\chi}_n(\omega, \phi_n)$, either the return to exiting, ϕ_n, or the return to remaining in the industry. The latter is given by the term within brackets and is in turn composed of two parts: the investment $\tilde{x}_n(\omega, \phi_n)$ and the net present value of the incumbent's future cash flows, $\beta \mathrm{E}\{V_n(\omega')|\cdot\}$. Several remarks are in order. First, because scrap values are independent across periods, the firm's future returns are described by its expected value function $V_n(\omega')$. Second, recall that ω' denotes the state at the end of the current period after the state-to-state transitions have been realized. The expectation operator reflects the fact that ω' is unknown at the beginning of the current period when the decisions are made. The incumbent conditions its expectations on the decisions of its rivals, $\xi_{-n}(\omega)$ and $x_{-n}(\omega)$. It also conditions its expectations on its own investment choice and presumes that it remains in the industry in state ω, that is, it conditions on $\omega'_n \neq M+1$.

Because investment is chosen conditional on remaining in the industry, the problem of incumbent firm n can be broken up into two parts. First, the incumbent chooses its investment. The optimal investment choice is independent of the firm's scrap value, and there is thus no need to index $x_n(\omega)$ by ϕ_n. This also justifies making the expectation operator conditional on $x_{-n}(\omega)$ (as opposed to scrap-value-specific investment decisions). Second, given its investment choice, the incumbent decides whether or not to remain in the industry. The incumbent's exit decision clearly depends on its scrap value, just as its rivals' exit and entry decisions depend on their scrap values and setup costs. Nevertheless, it is enough to condition on $\xi_{-n}(\omega)$ in light of equation (1).

The optimal exit decision of incumbent firm n who has drawn scrap value ϕ_n is a cutoff rule characterized by

$$\chi_n(\omega, \phi_n) = \begin{cases} 1 & \text{if } \phi_n < \bar{\phi}_n(\omega), \\ 0 & \text{if } \phi_n \geq \bar{\phi}_n(\omega), \end{cases}$$

where

$$\bar{\phi}_n(\omega) = \sup_{\tilde{x}_n(\omega) \in [0,\bar{x}]} -\tilde{x}_n(\omega) + \beta \mathrm{E}\{V_n(\omega')|\omega, \omega'_n \neq M+1, \tilde{x}_n(\omega), \xi_{-n}(\omega), x_{-n}(\omega)\} \qquad (3)$$

denotes the cutoff scrap value for which the incumbent is indifferent between remaining in the industry and exiting. Hence, the solution to the incumbent's decision problem has the reservation property. Moreover, under appropriate assumptions on $F(\cdot)$ (see Section 4), incumbent firm n has a unique optimal exit choice for all scrap values. Without loss of generality, we can therefore restrict attention to decision rules of the form $1[\phi_n < \bar{\phi}_n(\omega)]$, where $1[\cdot]$ denotes the indicator function. These decision rules can be represented in two ways:

(i) with the cutoff scrap value $\bar{\phi}_n(\omega)$ itself; or
(ii) with the probability $\xi_n(\omega)$ of incumbent firm n remaining in the industry in state ω.

This is without loss of information because $\xi_n(\omega) = \int \chi_n(\omega, \phi_n) dF(\phi_n) = \int 1[\phi_n < \bar{\phi}_n(\omega)] dF(\phi_n) = F(\bar{\phi}_n(\omega))$ is equivalent to $F^{-1}(\xi_n(\omega)) = \bar{\phi}_n(\omega)$.[12] The second representation proves to be more useful, and we use it below almost exclusively.

Next we turn to payoffs. Imposing the reservation property and integrating over ϕ_n on both sides of equation (2) yields

$$\begin{aligned} V_n(\omega) &= \int \sup_{\substack{\tilde{\xi}_n(\omega) \in [0,1], \\ \tilde{x}_n(\omega) \in [0,\bar{x}]}} \pi_n(\omega) + (1 - 1[\phi_n < F^{-1}(\tilde{\xi}_n(\omega))])\phi_n + 1[\phi_n < F^{-1}(\tilde{\xi}_n(\omega))] \\ &\quad \times \{-\tilde{x}_n(\omega) + \beta \mathrm{E}\{V_n(\omega') | \omega, \omega'_n \neq M+1, \tilde{x}_n(\omega), \xi_{-n}(\omega), x_{-n}(\omega)\}\} dF(\phi_n) \\ &= \sup_{\substack{\tilde{\xi}_n(\omega) \in [0,1], \\ \tilde{x}_n(\omega) \in [0,\bar{x}]}} \pi_n(\omega) + (1 - \tilde{\xi}_n(\omega))\phi + \int_{\phi_n > F^{-1}(\tilde{\xi}_n(\omega))} (\phi_n - \phi) dF(\phi_n) + \tilde{\xi}_n(\omega) \\ &\quad \times \{-\tilde{x}_n(\omega) + \beta \mathrm{E}\{V_n(\omega') | \omega, \omega'_n \neq M+1, \tilde{x}_n(\omega), \xi_{-n}(\omega), x_{-n}(\omega)\}\}. \end{aligned} \tag{4}$$

Two essential points should be noted. First, an optimizing incumbent cares about the expectation of the scrap value conditional on collecting it, $\mathrm{E}\{\phi_n | \phi_n > F^{-1}(\tilde{\xi}_n(\omega))\}$, rather than its unconditional expectation, $\mathrm{E}(\phi_n) = \phi$. The term $\int_{\phi_n > F^{-1}(\tilde{\xi}_n(\omega))} (\phi_n - \phi) dF(\phi_n) = (1 - \tilde{\xi}_n(\omega))(\mathrm{E}\{\phi_n | \phi_n > F^{-1}(\tilde{\xi}_n(\omega))\} - \phi)$ captures the difference between the conditional and the unconditional expectation. It reflects our assumption that scrap values are random and, consequently, it is not present in a game of complete information such as Ericson and Pakes (1995) where scrap values are constant across firms and periods. Second, the state space is effectively the same in the games of incomplete and complete information, because the constituent parts of the Bellman equation (4) depend on the state of the industry ω but not on the random scrap value ϕ_n. Hence, by integrating out over the random scrap values, we have successfully eliminated their computational disadvantage.

□ **An entrant's problem.** Suppose that the industry is in state ω with $\omega_n = M + 1$. The expected net present value of all future cash flows to potential entrant n when the industry is in state ω and the firm has drawn setup cost ϕ_n^e is

$$\begin{aligned} V_n(\omega, \phi_n^e) = \sup_{\substack{\tilde{\chi}_n(\omega, \phi_n^e) \in \{0,1\}, \\ \tilde{x}_n(\omega) \in [0,\bar{x}^e]}} &\tilde{\chi}_n(\omega, \phi_n^e) \big\{ -\phi_n^e - \tilde{x}_n(\omega) \\ &+ \beta \mathrm{E}\{V_n(\omega') | \omega, \omega'_n \neq M+1, \tilde{x}_n(\omega), \xi_{-n}(\omega), x_{-n}(\omega)\}\big\}. \end{aligned} \tag{5}$$

Because the entrant is short lived, it does not solve an intertemporal maximization problem to reach its decisions.[13] Depending on the entry decision $\chi_n(\omega, \phi^e)$, the right-hand side of the above equation is either 0 or the expected return to entering the industry, which is in turn composed of two parts. First, the entrant pays the setup cost and sinks its investment, yielding a current cash flow of $-\phi_n^e - \tilde{x}_n(\omega)$. Second, the entrant takes the net present value of its future cash flows into

[12] If the support of $F(\cdot)$ is bounded, we define $F^{-1}(0)$ to be its minimum and $F^{-1}(1)$ to be its maximum.

[13] It is easy to allow for long-lived entrants by adding the term $(1 - \tilde{\chi}_n(\omega, \phi_n^e))\beta \mathrm{E}\{V_n(\omega') | \omega, \omega'_n = M + 1, \tilde{x}_n(\omega, \phi_n^e), \xi_{-n}(\omega), x_{-n}(\omega)\}$ to equation (5).

account. Because potential entrant n becomes incumbent firm n at the end of the setup period, this is given by $\beta \mathrm{E}\{V_n(\omega')|\cdot\}$. The entrant conditions its expectations on the decisions of its rivals, $\xi_{-n}(\omega)$ and $x_{-n}(\omega)$. It also conditions its expectations on its own investment choice and presumes that it enters the industry in state ω, that is, it conditions on $\omega'_n \neq M+1$.

Similar to the incumbent's problem, the entrant's problem can be broken up into two parts. Because investment is chosen conditional on entering the industry, the optimal investment choice $x_n(\omega)$ is independent of the firm's setup cost ϕ^e_n. Given its investment choice, the entrant then decides whether or not to enter the industry. The optimal entry decision is characterized by

$$\chi_n(\omega, \phi^e_n) = \begin{cases} 1 & \text{if } \phi^e_n \leq \bar{\phi}^e_n(\omega), \\ 0 & \text{if } \phi^e_n > \bar{\phi}^e_n(\omega), \end{cases}$$

where

$$\bar{\phi}^e_n(\omega) = \sup_{\tilde{x}_n(\omega)\in[0,\bar{x}^e]} -\tilde{x}_n(\omega) + \beta \mathrm{E}\left\{V_n(\omega')|\omega, \omega'_n \neq M+1, \tilde{x}_n(\omega), \xi_{-n}(\omega), x_{-n}(\omega)\right\} \tag{6}$$

denotes the cutoff setup cost. As with incumbents, the solution to the entrant's decision problem has the reservation property and we can restrict attention to decision rules of the form $1[\phi^e_n < \bar{\phi}^e_n(\omega)]$ that can be alternatively represented by the probability $\xi_n(\omega)$ of potential entrant n entering the industry in state ω. Imposing the reservation property and integrating over ϕ^e_n on both sides of equation (5) yields

$$V_n(\omega) = \sup_{\substack{\tilde{\xi}_n(\omega)\in[0,1],\\ \tilde{x}_n(\omega)\in[0,\bar{x}^e]}} -\int_{\phi^e_n < F^{e-1}(\tilde{\xi}_n(\omega))} \left(\phi^e_n - \phi^e\right) dF^e\left(\phi^e_n\right) + \tilde{\xi}_n(\omega)$$
$$\times\{-\phi^e - \tilde{x}_n(\omega) + \beta \mathrm{E}\{V_n(\omega')|\omega, \omega'_n \neq M+1, \tilde{x}_n(\omega), \xi_{-n}(\omega), x_{-n}(\omega)\}\}. \tag{7}$$

The term $-\int_{\phi^e_n < F^{e-1}(\tilde{\xi}_n(\omega))}(\phi^e_n - \phi^e)dF^e(\phi^e_n)$ is again not present in a setting with complete information.

□ **Actions, strategies, and payoffs.** An action or decision for firm n in state ω specifies either the probability that the incumbent remains in the industry or the probability that the entrant enters the industry along with an investment choice: $u_n(\omega) = (\xi_n(\omega), x_n(\omega)) \in \mathcal{U}_n(\omega)$, where

$$\mathcal{U}_n(\omega) = \begin{cases} [0,1]\times[0,\bar{x}] & \text{if } \omega_n \neq M+1, \\ [0,1]\times[0,\bar{x}^e] & \text{if } \omega_n = M+1 \end{cases} \tag{8}$$

denotes firm n's feasible actions in state ω.

We restrict attention to stationary Markovian strategies. A strategy or policy for firm n is a single function from states into actions; it specifies an action $u_n(\omega) \in \mathcal{U}_n(\omega)$ for each state ω. Such a strategy is called Markovian because it is restricted to be a function of the current state rather than the entire history of the game. It is called stationary because it does not directly depend on calendar time, that is, the firm plays the same action $u_n(\omega)$ each time the industry is in state ω.[14]

Define $\mathcal{U}_n = \times_{\omega\in S}\mathcal{U}_n(\omega)$ to be the strategy space of firm n. Any element of the set $\mathcal{U}_n$ is a stationary Markovian strategy. Further define $\mathcal{U} = \times_{n=1}^N \mathcal{U}_n$ to be the strategy space of the entire industry. By construction in equation (8), the set of feasible actions $\mathcal{U}_n(\omega)$ is nonempty, convex, and compact (as long as $\bar{x} < \infty$ and $\bar{x}^e < \infty$). It follows that the strategy spaces $\mathcal{U}_n$ and $\mathcal{U}$ are also nonempty, convex, and compact.

Turning to payoffs, the Bellman equations (4) and (7) of incumbent firm n and potential entrant n, respectively, can be more compactly stated as

$$V_n(\omega) = \sup_{\tilde{u}_n\in\mathcal{U}_n(\omega)} h_n(\omega, \tilde{u}_n(\omega), u_{-n}(\omega), V_n), \tag{9}$$

[14] Nonstationary strategies are computationally infeasible in infinite-horizon models like ours because they require computing a different function for each period. Stationarity is also a compelling modeling restriction whenever nothing in the economic environment depends directly on calendar time.

where

$$h_n(\omega, u(\omega), V_n) = \begin{cases} \pi_n(\omega) + (1 - \xi_n(\omega))\phi + \int_{\phi_n > F^{-1}(\xi_n(\omega))} (\phi_n - \phi) dF(\phi_n) \\ + \xi_n(\omega)\Big\{ - x_n(\omega) + \beta \mathrm{E}\left\{ V_n(\omega') | \omega, \omega'_n \neq M+1, \xi_{-n}(\omega), x(\omega) \right\}\Big\} & \text{if } \omega_n \neq M+1, \\ - \int_{\phi_n^e < F^{e-1}(\xi_n(\omega))} \left(\phi_n^e - \phi^e\right) dF^e\left(\phi_n^e\right) \\ + \xi_n(\omega)\Big\{ - \phi^e - x_n(\omega) + \beta \mathrm{E}\left\{ V_n(\omega') | \omega, \omega'_n \neq M+1, \xi_{-n}(\omega), x(\omega) \right\}\Big\} & \text{if } \omega_n = M+1. \end{cases} \tag{10}$$

The number $h_n(\omega, u(\omega), V_n)$ represents the return to firm n in state ω when the firms use actions $u(\omega)$ and firm n's future returns are described by the value function V_n. The function $h_n(\cdot)$ is called firm n's return (Denardo, 1967) or local income function (Whitt, 1980).

Enumerate the state space as $S = \Omega^N = \{\omega^1, \ldots, \omega^{|S|}\}$ and define the $|S| \times N$ matrix V by

$$V = (V_1, \ldots, V_N) = \begin{pmatrix} V_1(\omega^1) & \ldots & V_N(\omega^1) \\ \vdots & & \vdots \\ V_1(\omega^{|S|}) & \ldots & V_N(\omega^{|S|}) \end{pmatrix}$$

and the $|S| \times (N-1)$ matrix V_{-n} by $V_{-n} = (V_1, \ldots, V_{n-1}, V_{n+1}, \ldots, V_N)$. V_n represents the value function of firm n or, more precisely, the value function of incumbent firm n if $\omega_n \neq M+1$ and the value function of potential entrant n if $\omega_n = M+1$. Define $V(\omega) = (V_1(\omega), \ldots, V_N(\omega))$ and $V_{-n}(\omega) = (V_1(\omega), \ldots, V_{n-1}(\omega), V_{n+1}(\omega), \ldots, V_N(\omega))$. Define the $|S| \times N$ matrices ξ and x similarly. Finally, define the $|S| \times 2N$ matrix u by $u = (\xi, x)$. In what follows, we use the terms matrix and function interchangeably.

□ **Equilibrium.** Our solution concept is that of stationary Markov-perfect equilibrium. An equilibrium involves value and policy functions V and u such that (i) given u_{-n}, V_n solves the Bellman equation (9) for all n and (ii) given $u_{-n}(\omega)$ and V_n, $u_n(\omega)$ solves the maximization problem on the right-hand side of this equation for all ω and all n. A firm thus behaves optimally in every state, irrespective of whether this state is on or off the equilibrium path. Moreover, because the horizon is infinite and the influence of past play is captured in the current state, there is a one-to-one correspondence between subgames and states. Hence, any stationary Markov-perfect equilibrium is subgame perfect. Note that because a best reply to stationary Markovian strategies u_{-n} is a stationary Markovian strategy u_n, a stationary Markov-perfect equilibrium remains a subgame-perfect equilibrium even if nonstationary strategies are considered. Of course, this does not rule out that there may also exist nonstationary Markov-perfect equilibria.

3. Examples

■ In this section, we provide two simple examples to illustrate the key themes of the subsequent analysis. Our first example demonstrates that if scrap values/setup costs are constant across firms and periods as in the Ericson and Pakes' (1995) model, then a symmetric equilibrium in pure entry/exit strategies may fail to exist, contrary to their assertion.[15] Our second example shows how to incorporate random scrap values/setup costs in order to ensure that a symmetric equilibrium in cutoff entry/exit strategies exists.

□ **Example: deterministic scrap values/setup costs.** We set $N = 2$ and $M = 1$. This implies that the industry is either a monopoly (states (1,2) and (2,1)) or a duopoly (state (1,1)). Moreover,

[15] We defer a formal definition of our symmetry notion to Section 6.

because there is just one "active" state, there is no incentive to invest, so we set $x_n(\omega) = 0$ for all ω and all n in what follows. To simplify further, we assume that entry is prohibitively costly and focus entirely on exit. Let $\pi(\omega_1, \omega_2)$ denote firm 1's current profit in state $\omega = (\omega_1, \omega_2)$. We assume that the profit function is symmetric. This implies that firm 2's current profit in state ω is $\pi(\omega_2, \omega_1)$. Pick the deterministic scrap value ϕ such that

$$\frac{\beta\pi(1,1)}{1-\beta} < \phi < \frac{\beta\pi(1,2)}{1-\beta}. \tag{11}$$

Hence, whereas a monopoly is viable, a duopoly is not. This gives rise to a "war of attrition."

The sole decision that a firm must make is whether or not to exit the industry. Consider firm 1. Given firm 2's exit decision $\chi(1,1) \in \{0,1\}$, the Bellman equation defines its value function:

$$V(1,2) = \sup_{\tilde{\chi}(1,2)\in\{0,1\}} \pi(1,2) + (1-\tilde{\chi}(1,2))\phi + \tilde{\chi}(1,2)\beta V(1,2),$$

$$V(1,1) = \sup_{\tilde{\chi}(1,1)\in\{0,1\}} \pi(1,1) + (1-\tilde{\chi}(1,1))\phi + \tilde{\chi}(1,1)\beta\{\chi(1,1)V(1,1) + (1-\chi(1,1))V(1,2)\}.$$

Recall that $\tilde{\chi}(\omega) = 1$ indicates that firm 1 remains in the industry in state ω and $\tilde{\chi}(\omega) = 0$ indicates that it exits. The optimal exit decisions $\tilde{\chi}(1,2)$ and $\tilde{\chi}(1,1)$ of firm 1 satisfy

$$\tilde{\chi}(\omega) = \begin{cases} 1 & \text{if } \phi \le \bar{\phi}(\omega), \\ 0 & \text{if } \phi \ge \bar{\phi}(\omega), \end{cases}$$

where

$$\bar{\phi}(1,2) = \beta V(1,2), \tag{12}$$

$$\bar{\phi}(1,1) = \beta\{\chi(1,1)V(1,1) + (1-\chi(1,1))V(1,2)\}. \tag{13}$$

Moreover, in a symmetric equilibrium, we must have $\tilde{\chi}(\omega_1, \omega_2) = \chi(\omega_2, \omega_1)$.

To show that there is no symmetric equilibrium in pure exit strategies, we show that $(\chi(1,2), \chi(1,1)) \in \{(0,0), (0,1), (1,0), (1,1)\}$ leads to a contradiction. Working through these cases, suppose first that $\chi(1,2) = 0$. Then $V(1,2) = \pi(1,2) + \phi$, and the assumed optimality of $\chi(1,2) = 0$ implies

$$\phi \ge \bar{\phi}(1,2) = \beta(\pi(1,2) + \phi) \Leftrightarrow \phi \ge \frac{\beta\pi(1,2)}{1-\beta}.$$

This contradicts assumption (11); therefore no equilibrium with $\chi(1,2) = 0$ exists. Next consider $\chi(1,1) = 1$. Then $V(1,1) = \frac{\pi(1,1)}{1-\beta}$, and the assumed optimality of $\chi(1,1) = 1$ implies

$$\phi \le \bar{\phi}(1,1) = \frac{\beta\pi(1,1)}{1-\beta}.$$

This contradicts assumption (11); therefore no equilibrium with $\chi(1,1) = 1$ exists. This leaves us with one more possibility: $\chi(1,2) = 1$ and $\chi(1,1) = 0$. Here $V(1,2) = \frac{\pi(1,2)}{1-\beta}$, and the assumed optimality of $\chi(1,2) = 1$ implies

$$\phi \ge \bar{\phi}(1,1) = \frac{\beta\pi(1,2)}{1-\beta},$$

which again contradicts assumption (11). Hence, there cannot be a symmetric equilibrium in pure exit strategies.[16]

[16] In this particular example, there exist two asymmetric equilibria in pure exit strategies. In each of them, within state (1,1), one firm exits for sure and the other firm remains for sure.

For future reference we note that, although there is no symmetric equilibrium in pure exit strategies, there is one in mixed exit strategies, given by

$$V(1,2) = \frac{\pi(1,2)}{1-\beta}, \quad V(1,1) = \pi(1,1) + \phi,$$

$$\xi(1,2) = 1, \quad \xi(1,1) = \frac{(1-\beta)\phi - \beta\pi(1,2)}{\beta\left((1-\beta)(\pi(1,1)+\phi) - \pi(1,2)\right)}.$$

□ **Example: random scrap values/setup costs.** Pakes and McGuire (1994) suggest the use of random setup costs to overcome convergence problems in their algorithm. Convergence problems may be indicative of nonexistence in pure entry/exit strategies. In the example above, an algorithm that seeks a (nonexistent) symmetric equilibrium in pure strategies tends to cycle between prescribing that neither firm should exit from a duopolistic industry and prescribing that both firms should exit.

To restore existence, we assume that scrap values are independently and identically distributed across firms and periods, and that its scrap value is private to itself. We write firm 1's scrap value as $\phi + \epsilon\theta$, where $\epsilon > 0$ is a constant scale factor that measures the importance of incomplete information. Overloading notation, we assume that $\theta \sim F(\cdot)$ with $\mathrm{E}(\theta) = 0$. The Bellman equation of firm 1 is

$$V(1,2) = \sup_{\tilde{\xi}(1,2)\in[0,1]} \pi(1,2) + (1-\tilde{\xi}(1,2))\phi + \epsilon \int_{\theta > F^{-1}(\tilde{\xi}(1,2))} \theta dF(\theta) + \tilde{\xi}(1,2)\beta V(1,2),$$

$$V(1,1) = \sup_{\tilde{\xi}(1,1)\in[0,1]} \pi(1,1) + (1-\tilde{\xi}(1,1))\phi + \epsilon \int_{\theta > F^{-1}(\tilde{\xi}(1,1))} \theta dF(\theta)$$
$$+ \tilde{\xi}(1,1)\beta\{\xi(1,1)V(1,1) + (1-\xi(1,1))V(1,2)\},$$

where $\xi(1,1) \in [0,1]$ is firm 2's exit decision. The optimal exit decisions of firm 1, $\tilde{\xi}(1,2)$ and $\tilde{\xi}(1,1)$, are characterized by $\tilde{\xi}(\omega) = F(\frac{\bar{\phi}(\omega)-\phi}{\epsilon})$, where[17]

$$\bar{\phi}(1,2) = \beta V(1,2),$$
$$\bar{\phi}(1,1) = \beta\{\xi(1,1)V(1,1) + (1-\xi(1,1))V(1,2)\}.$$

Moreover, in a symmetric equilibrium, we must have $\tilde{\xi}(\omega_1,\omega_2) = \xi(\omega_2,\omega_1)$. This yields a system of four equations in four unknowns: $V(1,2)$, $V(1,1)$, $\xi(1,2)$, and $\xi(1,1)$.

To facilitate the analysis, let θ be uniformly distributed on the interval $[-1,1]$.[18] This implies

$$\int_{\theta > F^{-1}(\xi(\omega))} \theta dF(\theta) = \begin{cases} 0 & \text{if } F^{-1}(\xi(\omega)) \le -1, \\ \dfrac{1 - F^{-1}(\xi(\omega))^2}{4} & \text{if } -1 < F^{-1}(\xi(\omega)) < 1, \\ 0 & \text{if } F^{-1}(\xi(\omega)) \ge 1, \end{cases}$$

where $F^{-1}(\xi(\omega)) = 2\xi(\omega) - 1$. There are nine cases to be considered, depending on whether $\xi(1,1)$ is equal to 0, between 0 and 1, or equal to 1 and on whether $\xi(1,2)$ is equal to 0, between 0 and 1, or equal to 1. Table 1 specifies parameter values.

A case-by-case analysis shows that, with random scrap values, there always exists a unique symmetric equilibrium. If $\epsilon > 5$, the equilibrium involves $0 < \xi(1,2) < 1$ and $0 < \xi(1,1) < 1$, and if $\epsilon \le 5$, it involves $\xi(1,2) = 1$ and $0 < \xi(1,1) < 1$. Table 2 describes the equilibrium for

[17] To see this, note that the first and second derivatives of the right-hand side of the Bellman equation are given by $\frac{d(.)}{d\tilde{\xi}(\omega)} = -\phi - \epsilon F^{-1}(\tilde{\xi}(\omega)) + \bar{\phi}(\omega)$ and $\frac{d^2(.)}{d\tilde{\xi}(\omega)^2} = -\epsilon \frac{1}{F'(F^{-1}(\tilde{\xi}(\omega)))}$, respectively.

[18] Besides the uniform distribution, many others yield a closed-form expression, including triangular and Beta distributions.

TABLE 1 Parameter Values

Parameter	$\pi(1,1)$	$\pi(1,2)$	ϕ	β
Value	0	1	15	$\frac{20}{21}$

TABLE 2 Equilibrium with Random Scrap Values

ϵ	$V(1,2)$	$V(1,1)$	$\xi(1,2)$	$\xi(1,1)$
10	23.817544	21.159671	0.884169	0.784836
5	21	18.044922	1	0.780375
2	21	16.392989	1	0.834562
1	21	15.730888	1	0.854920
0.1	21	15.076219	1	0.873034
0.01	21	15.007653	1	0.874804
0.001	21	15.000766	1	0.874980
10^{-6}	21	15.000001	1	0.875000

various values of ϵ. Given the parameter values in Table 1, the symmetric equilibrium in mixed strategies of the game of complete information is $V(1,2) = 21$, $V(1,1) = 15$, $\xi(1,2) = 1$, and $\xi(1,1) = \frac{7}{8} = 0.875$. As Table 2 shows, the equilibrium with random scrap values converges to the equilibrium in mixed strategies as ϵ approaches zero. In Sections 4 and 7, we show that existence and convergence are general properties of the game of incomplete information.

4. Existence

■ In this section, we show how incorporating firm heterogeneity in the form of random scrap values/setup costs into the Ericson and Pakes' (1995) model guarantees the existence of an equilibrium. We specifically establish the existence of a possibly asymmetric equilibrium. The proof extends Whitt (1980) to our setting. In fact, for the most part, it is a reassembly of his argument and some general results on dynamic programming due to Denardo (1967). Both articles use models that are sufficiently abstract to enable us to construct the bulk of the existence proof by citing their intermediate results. In developing our argument, we assume that firm n's investment problem always has a unique solution in order to guarantee that the equilibrium is in pure investment strategies. We state this assumption in terms of the local income function $h_n(\cdot)$. We then devote Section 5 to providing a sufficient condition in terms of the model's primitives for this assumption to hold.

We begin with a series of assumptions. The first one ensures that the model's primitives are bounded.

Assumption 1. (i) The state space is finite, that is, $N < \infty$ and $M < \infty$. (ii) Profits are bounded, that is, there exists $\bar{\pi} < \infty$ such that $-\bar{\pi} < \pi_n(\omega) < \bar{\pi}$ for all ω and all n. (iii) Investments are bounded, that is, $\bar{x} < \infty$ and $\bar{x}^e < \infty$. (iv) Scrap values and setup costs are drawn from distributions $F(\cdot)$ and $F^e(\cdot)$ that have both positive densities over connected supports and their expectations exist, that is, there exist $\bar{\phi} < \infty$ and $\bar{\phi}^e < \infty$ such that $-\bar{\phi} < \int |\phi_n| dF(\phi_n) < \bar{\phi}$ and $-\bar{\phi}^e < \int |\phi_n^e| dF^e(\phi_n^e) < \bar{\phi}^e$. (v) Firms discount future payoffs, that is, $\beta \in [0, 1)$.

The assumption in part (iii) is without loss of generality because the upper bounds $\bar{x}$ and $\bar{x}^e$ can always be chosen large enough to never constrain firms' optimal investment choices. Specifically, the best possible net present value of the current and future cash flows that any firm, be it an incumbent or an entrant, can realize is no greater than $\bar{V}^* = \bar{\phi}^e + \frac{\bar{\pi}}{1-\beta} + \bar{\phi}$, which is the sum of a bound on its entry subsidy (i.e., negative setup cost), the capitalized value of remaining in the best possible state forever, and a bound on its scrap value. Conversely, because a firm always has the option of investing zero, it can guarantee that the net present value of its current and future cash flows is no worse than $-\bar{V}^*$. Because no firm is ever willing to invest more than $\beta(\bar{V}^* - (-\bar{V}^*)) = 2\beta\bar{V}^*$ in order to reap the best instead of the worst possible net present value, upper bounds on investment in excess of $2\beta\bar{V}^*$ never constrain firms' optimal choices.

The assumption in part (iv) admits distributions $F(\cdot)$ and $F^e(\cdot)$ with either bounded or unbounded support. From a theorist's perspective, it is natural to assume bounded supports because unbounded supports essentially stipulate that some agent is willing to pay an arbitrarily large amount to acquire the assets of a firm that makes bounded profits from product market competition. From an empiricist's perspective, unbounded supports (as assumed by Aguirregabiria and Mira, 2007 and Pesendorfer and Schmidt-Dengler, 2008) may be attractive because they guarantee that in the data there cannot be an observation that has zero probability of occurring.

Next we assume continuity of the transition function $P(\cdot)$. Similar continuity assumptions are commonplace in the literature on dynamic stochastic games (see Mertens, 2002).

Assumption 2. $P(\omega', \omega, \chi(\omega, \phi), x(\omega))$ is a continuous function of $x(\omega)$ for all ω', ω, and all $\chi(\omega, \phi)$.

Observe from equation (10) that current profit is additively separable from investment. The continuity of the transition function $P(\cdot)$ in $x(\omega)$ therefore ensures the continuity of the local income function $h_n(\cdot)$ in $x(\omega)$. In addition, $h_n(\cdot)$ is continuous in $\xi(\omega)$ because, analogous to equation (1), firm n integrates out over all possible realizations of its rivals' exit and entry decisions $\chi_{-n}(\omega, \phi_{-n})$ to obtain the probability that the industry transits from state ω to state ω'. Observe further that $h_n(\cdot)$ is always continuous in V_n because V_n enters $h_n(\cdot)$ in equation (10) via the expected value of firm n's future cash flows, $\mathrm{E}\{V_n(\omega')|\cdot\}$. We record these observations for later use.

Proposition 1. Under Assumption 2, $h_n(\omega, u(\omega), V_n)$ is a continuous function of $u(\omega)$ and V_n for all ω and all n.

Due to the random scrap values/setup costs, our model is formally a dynamic stochastic game with a finite state space and a continuum of actions given by the probability that an incumbent firm remains in the industry/a potential entrant enters the industry and the set of feasible investment choices. Under Assumptions 1 and 2, standard arguments (e.g., Federgruen, 1978; Whitt, 1980) yield the existence of an equilibrium in mixed strategies. However, mixed strategies over continuous actions are infeasible to compute. To guarantee the existence of an equilibrium in pure investment strategies, we make the additional assumption that firm n's investment problem always has a unique solution.

Assumption 3. A unique $x_n(\omega)$ exists that attains the maximum of $h_n(\omega, 1, x_n(\omega), u_{-n}(\omega), V_n)$ for all $u_{-n}(\omega)$, V_n, ω, and all n.[19]

In Section 5, we provide a sufficient condition for Assumption 3 to hold in terms of the model's primitives. Specifically, we define UIC admissibility of the transition function $P(\cdot)$ and prove that this condition ensures uniqueness of investment choice. Therefore, Assumption 3 holds and an equilibrium that is amenable to computation exists. Constructing our argument in this modular form makes it simple and transparent for other researchers to insert alternative sufficient conditions for uniqueness of investment choice into our proof and immediately obtain existence.

Recall that we assume entry and exit decisions are implemented before investment decisions are carried out. Thus, firm n chooses $x_n(\omega)$ to maximize $h_n(\omega, 1, x_n(\omega), u_{-n}(\omega), V_n)$ in accordance with equations (3) and (6), and the resulting investment choice also maximizes

[19] Assumption 3 can be weakened to hold for all possible maximal return functions $V^*_{n,u_{-n}} \in [-\bar{V}^*, \bar{V}^*]^{|S|}$.

$h_n(\omega, \xi_n(\omega), x_n(\omega), u_{-n}(\omega), V_n)$ for all $\xi_n(\omega) > 0$, $u_{-n}(\omega)$, V_n, ω, and all n. Clearly any investment would be optimal whenever an incumbent firm exits for sure or a potential entrant stays out for sure. Consequently, we adopt the following convention: if $\xi_n(\omega) = 0$, then we take $x_n(\omega)$ to have the value alluded to in Assumption 3. It follows that $h_n(\omega, \xi_n(\omega), x_n(\omega), u_{-n}(\omega), V_n)$ attains its maximum for a unique value of $x_n(\omega)$ independent of the value of $\xi_n(\omega)$. This is a natural convention because if there were even the slightest chance that firm n would remain in the industry even though it sets $\xi_n(\omega) = 0$, then the firm would want to choose this value of $x_n(\omega)$ as its investment.

The above assumptions ensure existence of an equilibrium.

Proposition 2. Under Assumptions 1, 2, and 3, an equilibrium exists in cutoff entry/exit and pure investment strategies.

The proof is based on the following idea.[20] Fix strategies u_{-n} and consider firm n's problem. Because its competitors' strategies are fixed, firm n has to solve a decision problem (as opposed to a game problem). We can thus employ dynamic programming techniques to analyze the firm's problem. In particular, a contraction mapping argument establishes that the firm's best reply to its competitors' strategies is well defined. It remains to show that there exists a fixed point in the firms' best-reply correspondences.

Before stating the proof of Proposition 2, we introduce and discuss a number of constructs. We start with the decision problem. Let $\mathcal{V}_n$ denote the space of bounded $|S| \times 1$ vectors endowed with the sup norm. Fix $u_{-n} \in \mathcal{U}_{-n}$ and define the *maximal return operator* $H^*_{n,u_{-n}} : \mathcal{V}_n \to \mathcal{V}_n$ pointwise by

$$(H^*_{n,u_{-n}} V_n)(\omega) = \sup_{\tilde{u}_n(\omega)\in\mathcal{U}_n(\omega)} h_n(\omega, \tilde{u}_n(\omega), u_{-n}(\omega), V_n).$$

The number $(H^*_{u_{-n}} V_n)(\omega)$ represents the return to firm n in state ω when firm n chooses its optimal action while the other firms use actions $u_{-n}(\omega)$ and firm n's future returns are described by V_n. Note that the right-hand side of the above equation coincides with the right-hand side of the Bellman equation (9).

Because profits and investments are bounded and the expectations of scrap values and setup costs exist by Assumption 1, $H^*_{n,u_{-n}}$ takes bounded vectors into bounded vectors. Application of Blackwell's sufficient conditions (Blackwell, 1965, Theorem 5; see also Stokey and Lucas, 1989, Theorem 3.3) shows that $H^*_{n,u_{-n}}$ is a contraction with modulus β. First, inspection of equation (10) shows that $V_n(\omega) \geq \hat{V}_n(\omega)$ for all ω implies $(H^*_{n,u-n} V_n)(\omega) \geq (H^*_{n,u_{-n}} \hat{V}_n)(\omega)$ for all ω ("monotonicity"). Second, given a constant $c \geq 0$, $(H^*_{n,u-n}(V_n + c))(\omega) \leq (H^*_{n,u-n} V_n)(\omega) + \beta c$ for all ω ("discounting").

Because $H^*_{n,u_{-n}}$ is a contraction, the contraction mapping theorem (see Stokey and Lucas, 1989, Theorem 3.2) implies that there exists a unique $V^*_{n,u_{-n}} \in \mathcal{V}_n$ that satisfies $V^*_{n,u_{-n}} = H^*_{n,u_{-n}} V^*_{n,u_{-n}}$ or, equivalently,

$$V^*_{n,u_{-n}}(\omega) = \sup_{\tilde{u}_n(\omega)\in\mathcal{U}_n(\omega)} h_n(\omega, \tilde{u}_n(\omega), u_{-n}(\omega), V^*_{n,u_{-n}}) \tag{14}$$

for all ω. The fixed point $V^*_{n,u_{-n}}$ of $H^*_{n,u_{-n}}$ is called the *maximal return function* given policies u_{-n}; it should be thought of as a mapping from $\mathcal{U}_{-n}$ into $\mathcal{V}_n$. Clearly, given u_{-n}, the maximal return function $V^*_{n,u_{-n}}$ solves the Bellman equation (9); it plays a major role in our existence proof.

Before proceeding to the existence proof, we introduce and discuss another operator. Fix $u \in \mathcal{U}$ and define the *return operator* $H_{n,u} : \mathcal{V}_n \to \mathcal{V}_n$ pointwise by

$$(H_{n,u} V_n)(\omega) = h_n(\omega, u(\omega), V_n).$$

The number $(H_u V_n)(\omega)$ represents the return to firm n in state ω when the firms use actions $u(\omega)$ and V_n describes firm n's future returns. Like $H^*_{n,u_{-n}}$, $H_{n,u}$ is a contraction with modulus β

[20] Given that standard arguments establish the existence of an equilibrium in mixed strategies, it actually suffices to show that a firm is never willing to mix. The reason that we start from first principles is that we need the machinery from the proof of Proposition 2 for the proofs of Propositions 3, 5, and 6.

that takes bounded vectors into bounded vectors. Hence, a unique $V_{n,u} \in \mathcal{V}_n$ exists that satisfies $V_{n,u} = H_{n,u} V_{n,u}$, that is,

$$V_{n,u}(\omega) = h_n(\omega, u(\omega), V_{n,u}) \tag{15}$$

for all ω. The fixed point $V_{n,u}$ of $H_{n,u}$ is called the *return function* given policies u; it should be thought of as a mapping from $\mathcal{U}$ into $\mathcal{V}_n$.

The return function $V_{n,u}$ and the maximal return function $V^*_{n,u_{-n}}$ are tightly connected. Because the return operator $H_{n,u}$ is monotonic (meaning that $V_n(\omega) \geq \hat{V}_n(\omega)$ for all ω implies $(H_{n,u} V_n)(\omega) \geq (H_{n,u} \hat{V}_n)(\omega)$ for all ω), Denardo (1967) establishes that

$$V^*_{n,u_{-n}}(\omega) = \sup_{\tilde{u}_n \in \mathcal{U}_n} V_{n,\tilde{u}_n,u_{-n}}(\omega) \tag{16}$$

for all ω, where $V_{n,\tilde{u}_n,u_{-n}}$ is the fixed point of the return operator given policy $(\tilde{u}_n, u_{-n})$.

With this machinery in place, we turn to the game problem. Consider the mapping $\Upsilon_n : \mathcal{U}_{-n} \rightarrow \mathcal{U}_n$ defined by

$$\Upsilon_n(u_{-n}) = \left\{ \tilde{u}_n \in \mathcal{U}_n : \tilde{u}_n(\omega) \in \arg \sup_{\tilde{u}_n(\omega) \in \mathcal{U}_n(\omega)} h_n(\omega, \tilde{u}_n(\omega), u_{-n}(\omega), V^*_{n,u_{-n}}) \text{ for all } \omega \right\}. \tag{17}$$

$\Upsilon_n(\cdot)$ is the best-reply correspondence of firm n and $\Upsilon_n(u_{-n})$ is the set of best replies of firm n given rivals' policies u_{-n}. Consider further the mapping $\Upsilon : \mathcal{U} \rightarrow \mathcal{U}$ obtained by stacking these best-reply correspondences. $\Upsilon(u) = (\Upsilon_1(u_{-1}), \ldots, \Upsilon_N(u_{-N}))$ is the set of best replies of firm 1 given rivals' policies u_{-1}, those of firm 2 given rivals' policies u_{-2}, etc. An equilibrium exists if there is a $u \in \mathcal{U}$ such that $u \in \Upsilon(u)$. To show that such a u exists, we show that $\Upsilon(\cdot)$ is, in fact, a continuous function to which Brouwer's fixed-point theorem applies.

Proof of Proposition 2. We begin by establishing that $\Upsilon(\cdot)$ is nonempty and upper hemicontinuous. Given policies u_{-n}, firm n's maximal return function $V^*_{n,u_{-n}}$ is well defined due to Assumption 1, as shown above. Fix ω. Proposition 1 states that firm n's local income function $h_n(\omega, u(\omega), V_n)$ is continuous in $u(\omega)$ and V_n. The maximand, $h_n(\omega, u_n(\omega), u_{-n}(\omega), V^*_{n,u_{-n}})$, in the definition of $\Upsilon_n(\cdot)$ in equation (17) is therefore continuous in $u_n(\omega)$ and u_{-n} if firm n's maximal return function $V^*_{n,u_{-n}}$ is continuous in u_{-n}. That this is so is established through appeal to two lemmas by Whitt (1980).

In Lemma 3.2 he states that if $H_{n,u} V_n$ is continuous in u for all V_n, then the return function $V_{n,u}$ is continuous in u.[21] This establishes that $V_{n,u}$ is a continuous function of u. In Lemma 3.1 he states that if $\mathcal{U}_n(\omega)$, firm n's set of feasible actions in state ω, is a compact metric space for all ω, if the state space S is countable, and if the return function $V_{n,u}$ is continuous in u, then $\sup_{\tilde{u}_n \in \mathcal{U}_n} V_{n,\tilde{u}_n,u_{-n}}(\omega)$ is continuous in u_{-n} for all ω. These requirements are satisfied. Equation (16) thus implies that $V^*_{n,u_{-n}}(\omega)$ is continuous in u_{-n} for all ω. This, of course, implies that firm n's maximal return function $V^*_{n,u_{-n}}$ is continuous in u_{-n}.

Because $h_n(\omega, u_n(\omega), u_{-n}(\omega), V^*_{n,u_{-n}})$ is continuous in $u_n(\omega)$ and u_{-n} and $\mathcal{U}_n(\omega)$ is compact and independent of u_{-n}, the theorem of the maximum (see Stokey and Lucas, 1989, Theorem 3.6) implies that $\arg \sup_{\tilde{u}_n(\omega) \in \mathcal{U}_n(\omega)} h_n(\omega, \tilde{u}_n(\omega), u_{-n}(\omega), V^*_{n,u_{-n}})$ is nonempty and upper hemicontinuous in u_{-n}. Because ω was arbitrary, this establishes that $\Upsilon_n(\cdot)$ is a nonempty and upper hemicontinuous correspondence that maps $\mathcal{U}_{-n}$ into $\mathcal{U}_n$. Hence, $\Upsilon(\cdot)$ is a nonempty and upper hemicontinuous correspondence that maps $\mathcal{U}$ into $\mathcal{U}$.

We next show that $\Upsilon(\cdot)$ is single valued. Recall that, given policies u_{-n}, firm n's maximal return function $V^*_{n,u_{-n}}$ is well defined, and consider firm n's best reply in state ω. Uniqueness of the investment choice follows from Assumption 3 and our convention covering the special case of $\xi_n(\omega) = 0$. This, in turn, implies that equations (3) and (6) give unique exit and entry

[21] It is immediate to verify that the return operator $H_{n,u}$ satisfies the boundedness, monotonicity, and contraction assumptions in Whitt (1980). Whitt (1980) denotes the return function $V_{n,u}$ by $v_\delta(\cdot, i)$ and the maximal return function $V^*_{n,u_{-n}}$ by $f_\delta(\cdot, i)$. We set $\mathcal{W}_n = \mathcal{V}_n$ to obtain a special case of Lemma 3.2 in Whitt (1980).

cutoffs, $\bar{\phi}_n(\omega)$ and $\bar{\phi}^e_n(\omega)$. Given that these cutoffs are unique, the corresponding exit and entry probabilities, $\xi_n(\omega) = F(\bar{\phi}_n(\omega))$ (if $\omega_n \neq M+1$) and $\xi_n(\omega) = F^e(\bar{\phi}^e_n(\omega))$ (if $\omega_n = M+1$), must be unique. Because ω was arbitrary, this establishes that $\Upsilon_n(\cdot)$ and hence $\Upsilon(\cdot)$ is single valued.

Because $\Upsilon(\cdot)$ is nonempty, single valued, and upper hemicontinuous, it is, in fact, a continuous function that maps the nonempty, convex, and compact set $\mathcal{U}$ into itself. Brouwer's fixed-point theorem therefore applies: a $u \in \mathcal{U}$ exists such that $u \in \Upsilon(u)$. *Q.E.D.*

5. A sufficient condition for pure investment strategies

■ Assumption 3 requires that the local income function $h_n(\omega, 1, x_n(\omega), u_{-n}(\omega), V_n)$ is maximized at a unique investment choice $x_n(\omega)$ for all $u_{-n}(\omega)$, V_n, ω, and all n. Fortunately, a judicious choice of transition probabilities guarantees that the investment choice is indeed unique. In this section, we first define UIC admissibility of the transition function $P(\cdot)$ and show in Proposition 3 that if this condition on the model's primitives is satisfied, then Assumption 3 holds. We then provide a series of examples of transition functions that are UIC admissible and provide a reasonable amount of flexibility.

Condition 1. The transition function $P(\cdot)$ is unique investment choice (UIC) admissible if, for all $\chi_{-n}(\omega, \phi_{-n})$, $x(\omega)$, ω', ω, and all n, the probability $P(\omega', \omega, 1, \chi_{-n}(\omega, \phi_{-n}), x(\omega))$ that the industry moves from state ω to state ω' given that firm n remains in the industry (or enters the industry if firm n is an entrant rather than an incumbent) can be written in a separable form as

$$K_n\left(\omega', \omega, \chi_{-n}\left(\omega, \phi_{-n}\right), x_{-n}(\omega)\right) Q_n(\omega, x_n(\omega)) + L_n\left(\omega', \omega, \chi_{-n}\left(\omega, \phi_{-n}\right), x_{-n}(\omega)\right), \quad (18)$$

where $Q_n(\omega, x(\omega))$ is twice differentiable, strictly increasing, and strictly concave in $x_n(\omega)$, that is,

$$\frac{d}{dx_n(\omega)} Q_n(\omega, x_n(\omega)) > 0, \qquad \frac{d^2}{dx_n(\omega)^2} Q_n(\omega, x_n(\omega)) < 0 \quad (19)$$

for all $x_n(\omega) \in [0, \bar{x}]$ (or $x_n(\omega) \in [0, \bar{x}^e]$ if firm n is an entrant rather than an incumbent).[22]

UIC admissibility ensures that firm n's local income function $h_n(\ldots, 1, x_n(\omega), \ldots)$ either is strictly concave—and therefore has a unique maximizer—in the interval $[0, \bar{x}]$ (or in the interval $[0, \bar{x}^e]$ if firm n is an entrant rather than an incumbent) or that the unique maximizer is a corner solution. This establishes the following.

Proposition 3. If the transition function $P(\cdot)$ is UIC admissible (Condition 1), then Assumption 3 holds.

Proof. Because the proof for a potential entrant is the same with $\bar{x}^e$ replacing $\bar{x}$, we focus on the investment problem of an incumbent firm in what follows. UIC admissibility ensures that the expected value of firm n's future cash flow, $\mathrm{E}\{V_n(\omega')|\omega, \omega'_n \neq M+1, \xi_{-n}(\omega), x(\omega)\}$, in its local income function $h_n(\ldots, 1, x_n(\omega), \ldots)$ can be written in a separable form as

$$A_n(\omega, u_{-n}(\omega), V_n) Q_n(\omega, x_n(\omega)) + B_n(\omega, u_{-n}(\omega), V_n). \quad (20)$$

To see this, recall from equation (1) that firm n has to integrate out over all possible realizations of its rivals' exit and entry decisions to obtain the probability that the industry moves from state ω to state ω'. Hence,

[22] Condition 1 can be generalized to allow for $Q(\cdot)$ to depend on $x_{-n}(\omega)$.

$$\sum_{\omega' \in S} V_n(\omega') \sum_{\iota_{-n} \in \{0,1\}^{N-1}} P(\omega', \omega, 1, \iota_{-n}, x(\omega)) \prod_{k \neq n} \xi_k(\omega)^{\iota_k} (1 - \xi_k(\omega))^{1-\iota_k}$$

$$= \sum_{\omega' \in S} V_n(\omega') \sum_{\iota_{-n} \in \{0,1\}^{N-1}} [K_n(\omega', \omega, \iota_{-n}, x_{-n}(\omega)) Q_n(\omega, x_n(\omega)) + L_n(\omega', \omega, \iota_{-n}, x_{-n}(\omega))]$$

$$\times \prod_{k \neq n} \xi_k(\omega)^{\iota_k} (1 - \xi_k(\omega))^{1-\iota_k}$$

$$= \underbrace{\left[\sum_{\omega' \in S} V_n(\omega') \sum_{\iota_{-n} \in \{0,1\}^{N-1}} K_n(\omega', \omega, \iota_{-n}, x_{-n}(\omega)) \prod_{k \neq n} \xi_k(\omega)^{\iota_k} (1 - \xi_k(\omega))^{1-\iota_k} \right]}_{A_n(\omega, u_{-n}(\omega), V_n)} Q_n(\omega, x_n(\omega))$$

$$+ \underbrace{\left[\sum_{\omega' \in S} V_n(\omega') \sum_{\iota_{-n} \in \{0,1\}^{N-1}} L_n(\omega', \omega, \iota_{-n}, x_{-n}(\omega)) \prod_{k \neq n} \xi_k(\omega)^{\iota_k} (1 - \xi_k(\omega))^{1-\iota_k} \right]}_{B_n(\omega, u_{-n}(\omega), V_n)},$$

where the first equality uses the separability condition (18).

Next we differentiate $h_n(\ldots, 1, x_n(\omega), \ldots)$ with respect to $x_n(\omega)$. By virtue of equation (20), the first-order condition for an unconstrained solution to firm n's investment problem is

$$-1 + \beta A_n(\omega, u_{-n}(\omega), V_n) \frac{d}{dx_n(\omega)} Q_n(\omega, x_n(\omega)) = 0.$$

There are two cases to consider. First suppose that $A_n(\omega, u_{-n}(\omega), V_n) > 0$. If there exists a solution to the first-order condition in $[0, \bar{x}]$, say $\hat{x}_n(\omega)$, then it must be unique because the objective function is strictly concave on $[0, \bar{x}]$ in light of the derivative condition (19). Hence, $x_n(\omega) = \hat{x}_n(\omega)$ is the unique maximizer. If there does not exist a solution to the first-order condition in $[0, \bar{x}]$, then the objective function is either strictly decreasing or strictly increasing on $[0, \bar{x}]$. In the former case the unique maximizer is $x_n(\omega) = 0$, and in the latter case it is $x_n(\omega) = \bar{x}$.

Next suppose that $A_n(\omega, u_{-n}(\omega), V_n) \leq 0$. The objective function is strictly decreasing. Hence, the unique maximizer is $x_n(\omega) = 0$. *Q.E.D.*

UIC admissibility allows for much more flexibility in the transition probabilities than the simple schemes seen in the extant literature where each firm each period is restricted to move up one state, stay the same, or drop down one state. We demonstrate this with a series of increasingly complex examples all involving an industry with $N = 2$ firms, $M \geq 3$ "active" states, and no entry and exit.

□ **Example: independent transitions to immediately adjacent states.** Consider a game of capacity accumulation similar to that in Besanko and Doraszelski (2004). A firm's state describes its capacity. In each period, the firm decides how much to spend on an investment project in order to add to its capacity. If firm n invests $x_n(\omega) \geq 0$, then the probability that its investment project succeeds is

$$p_n = \frac{\alpha x_n(\omega)}{1 + \alpha x_n(\omega)},$$

where the parameter $\alpha > 0$ measures the effectiveness of investment. Depreciation tends to offset investment, and we assume that each firm is independently hit by a depreciation shock with probability δ. The transition probabilities at an interior state $\omega \in \{2, \ldots, M-1\}^2$ are given in Table 3.

Without loss of generality, consider firm 1. The probability of remaining in state ω can be written as

TABLE 3 Transition Probabilities for Independent Transitions to Immediately Adjacent States

	$\omega_2' = \omega_2 + 1$	$\omega_2' = \omega_2$	$\omega_2' = \omega_2 - 1$
$\omega_1' = \omega_1 + 1$	$(1-\delta)p_1(1-\delta)p_2$	$(1-\delta)p_1[\delta p_2 + (1-\delta)(1-p_2)]$	$(1-\delta)p_1\delta p_2$
$\omega_1' = \omega_1$	$[\delta p_1 + (1-\delta)(1-p_1)]$ $\times (1-\delta)p_2$	$[\delta p_1 + (1-\delta)(1-p_1)]$ $\times [\delta p_2 + (1-\delta)(1-p_2)]$	$[\delta p_1 + (1-\delta)(1-p_1)]$ $\times \delta p_2$
$\omega_1' = \omega_1 - 1$	$\delta(1-p_1)(1-\delta)p_2$	$\delta(1-p_1)[\delta p_2 + (1-\delta)(1-p_2)]$	$\delta(1-p_1)\delta p_2$

TABLE 4 Transition Probabilities for Dependent Transitions to Immediately Adjacent States

	$\omega_2' = \omega_2 + 1$	$\omega_2' = \omega_2$	$\omega_2' = \omega_2 - 1$
$\omega_1' = \omega_1 + 1$	$(1-\delta)p_1 p_2$	$(1-\delta)p_1(1-p_2)$	0
$\omega_1' = \omega_1$	$(1-\delta)(1-p_1)p_2$	$(1-\delta)(1-p_1)(1-p_2) + \delta p_1 p_2$	$\delta p_1(1-p_2)$
$\omega_1' = \omega_1 - 1$	0	$\delta(1-p_1)p_2$	$\delta(1-p_1)(1-p_2)$

$$[\delta p_1 + (1-\delta)(1-p_1)][\delta p_2 + (1-\delta)(1-p_2)] = \underbrace{[2\delta - 1][\delta p_2 + (1-\delta)(1-p_2)]}_{K_1(\omega,\omega,x_2(\omega))} \underbrace{p_1}_{Q_1(\omega,x_1(\omega))} + \underbrace{[1-\delta][\delta p_2 + (1-\delta)(1-p_2)]}_{L_1(\omega,\omega,x_2(\omega))}.$$

This expression satisfies the separability condition (18), as do the corresponding expressions for the probabilities of moving to some other state $\omega' \neq \omega$. In addition, the derivative condition (19) is satisfied because

$$\frac{d}{dx_1(\omega)} Q_1(\omega, x_1(\omega)) = \frac{\alpha}{(1+\alpha x_1(\omega))^2} > 0, \quad \frac{d^2}{dx_1(\omega)^2} Q_1(\omega, x_1(\omega)) = -\frac{2\alpha^2}{(1+\alpha x_1(\omega))^3} < 0.$$

□ **Example: dependent transitions to immediately adjacent states.** Next we introduce correlation into firms' transitions by replacing the firm-specific depreciation shocks of the above example by an industry-wide depreciation shock (e.g., Pakes and McGuire, 1994). Decompose, for purposes of exposition, the transition of each firm into two stages. In the first stage, the probability that firm n's state increases by one is again given by p_n. In the second stage, a depreciation shock reduces the states of all firms by one with probability δ. The transition probabilities at an interior state $\omega \in \{2, \ldots, M-1\}^2$ are given in Table 4.

For the sake of brevity, we just spell out the probability of remaining in state ω,

$$(1-\delta)(1-p_1)(1-p_2) + \delta p_1 p_2 = \underbrace{[\delta - 1 + p_2]}_{K_1(\omega,\omega,x_2(\omega))} \underbrace{p_1}_{Q_1(\omega,x_1(\omega))} + \underbrace{[(1-\delta)(1-p_2)]}_{L_1(\omega,\omega,x_2(\omega))},$$

and note that conditions (18) and (19) are again both satisfied.

□ **Example: dependent transitions to arbitrary states.** Using the above two-stage decomposition, much more flexible transitions can be constructed. In the first stage, firm n's investment $x_n(\omega)$ determines a set of transition probabilities to all possible "active" states. For example, the probability that firm n moves from its initial state ω_n to the intermediate state $\hat{\omega}_n \in \{1, \ldots, M\}$ may be

$$\begin{cases} \zeta_{n,\omega_n,1} + \eta_{n,\omega_n,1}\, p_n & \text{if} \quad \hat{\omega}_n = 1, \\ \vdots & \vdots \quad \vdots \\ \zeta_{n,\omega_n,\omega_n-1} + \eta_{n,\omega_n,\omega_n-1}\, p_n & \text{if } \hat{\omega}_n = \omega_n - 1, \\ \zeta_{n,\omega_n,\omega_n} + \eta_{n,\omega_n,\omega_n}\, p_n & \text{if} \quad \hat{\omega}_n = \omega_n, \\ \zeta_{n,\omega_n,\omega_n+1} + \eta_{n,\omega_n,\omega_n+1}\, p_n & \text{if } \hat{\omega}_n = \omega_n + 1, \\ \vdots & \vdots \quad \vdots \\ \zeta_{n,\omega_n,M} + \eta_{n,\omega_n,M}\, p_n & \text{if} \quad \hat{\omega}_n = M, \end{cases}$$

where $x_n(\omega)$ affects the probability of a transition from state ω_n to state $\hat{\omega}_n$ either positively of negatively depending on the sign of $\eta_{n,\omega_n,\hat{\omega}_n}$.[23] In the second stage, the industry transits from its intermediate state $\hat{\omega}$ to its final state ω' according to some arbitrary, exogenously given probabilities that may depend on $\hat{\omega}$.

Clearly, p_n does not have to equal $\frac{\alpha x_n(\omega)}{1+\alpha x_n(\omega)}$; it can be of any form that satisfies the derivative condition (19). For example, let

$$p_n = 1 - e^{-\alpha x_n(\omega)},$$

where $\alpha > 0$. As another example, let

$$p_n = \frac{\arctan\left(\frac{2\alpha_1 x_n(\omega) + \alpha_2}{\sqrt{4-\alpha_2^2}}\right) - \arctan\left(\frac{\alpha_2}{\sqrt{4-\alpha_2^2}}\right)}{\frac{\pi}{2} - \arctan\left(\frac{\alpha_2}{\sqrt{4-\alpha_2^2}}\right)},$$

where $\alpha_1 > 0$ and $0 \leq \alpha_2 < 2$. Then p_n is increasing in α_1 (just as $\frac{\alpha x_n(\omega)}{1+\alpha x_n(\omega)}$ and $1 - e^{\alpha x_n(\omega)}$ are increasing in α) and increasing (decreasing) in α_2 to the left (right) of $x_n(\omega) = \frac{1}{\alpha_1}$. That is, whereas increasing α_1 makes investments of all sizes more effective, increasing α_2 makes small investments more and large ones less effective. In addition, $x_n(\omega) = \frac{1}{\alpha_1}$ implies $p_n = \frac{1}{2}$. Hence, increasing α_2 preserves the median but increases the spread of p_n as measured, for example, by the interquartile range.

UIC admissibility is a sufficient condition and, if it fails, uniqueness of investment choice can often be achieved by other means. Purification is again a very valuable tool. In particular, a part of the subsequent literature has assumed that the cost of investment is randomly drawn and privately known. Ryan (2006) and Besanko, Doraszelski, Lu, and Satterthwaite (2010) extend our handling of entry and exit to the case of discrete (or "lumpy") investment. Their models remains computationally tractable because the equilibrium is in cutoff investment strategies. Focusing on the case of continuous investment, Jenkins, Liu, Matzkin, and McFadden (2004) restrict the functional form of per-period payoffs to ensure that a firm's optimal investment level is almost always unique given a realization of the cost of investment. Again its rivals perceive the firm *as though* it were following a mixed strategy, thereby facilitating the existence of an equilibrium, although computing these perceptions—as one must in order to determine the rivals' best replies to them—becomes somewhat more involved.

6. Symmetry

■ In Section 4, we established the existence of a possibly asymmetric equilibrium. We now show that if the model's primitives satisfy an additional symmetry assumption, then a symmetric equilibrium exists.

[23] The parameters $\zeta_{n,\omega_n,\hat{\omega}_n}$ and $\eta_{n,\omega_n,\hat{\omega}_n}$ must be chosen to ensure that the probabilities stay in the unit interval for all $x_n(\omega) \in [0, \bar{x}]$ and sum to one. In particular, this requires $\sum_{\hat{\omega}_n=1}^{M} \zeta_{n,\omega_n,\hat{\omega}_n} = 1$ and $\sum_{\hat{\omega}_n=1}^{M} \eta_{n,\omega_n,\hat{\omega}_n} = 0$.

Informally, the notion of symmetry in Ericson and Pakes (1995) is this: consider an industry with five firms and suppose that when firm 2 is in state 3 and the other four firms are in states 1, 3, 3, and 6, it invests 50. Symmetry means that when firm 4 is in state 3 and the other four firms are in states 1, 3, 3, and 6, it also invests 50. Thus, in a symmetric equilibrium, a firm's policy is a *common* function of its *own* state and the *distribution* of its rivals' states.

To formalize this notion of symmetry, let $\kappa = (\kappa_1, \ldots, \kappa_N)$ be a permutation of $(1, \ldots, N)$. The policy functions $u = (u_1, \ldots, u_N)$ are symmetric if

$$u_n(\omega_{\kappa_1}, \ldots, \omega_{\kappa_{n-1}}, \omega_{\kappa_n}, \omega_{\kappa_{n+1}}, \ldots, \omega_{\kappa_N}) = u_{\kappa_n}(\omega_1, \ldots, \omega_{n-1}, \omega_n, \omega_{n+1}, \ldots, \omega_N) \qquad (21)$$

for all ω, n, and all κ. We say that an equilibrium is symmetric if its policy functions are symmetric. Moreover, in a symmetric equilibrium, the value functions $V = (V_1, \ldots, V_N)$ are symmetric and satisfy the analog of equation (21).

This definition implies two key properties that capture the essence of symmetry:

(i) If the states of two firms are the same, then their actions must be the same. For example, if $\omega = (2, 3, 3)$, then set $\kappa = (2, 3, 1)$ in equation (21) to obtain $u_2(3, 3, 2) = u_2(\omega_2, \omega_3, \omega_1) = u_3(\omega_1, \omega_2, \omega_3) = u_3(2, 3, 3)$.

(ii) A firm does not care about the identity of its rivals; hence, the firm's action must be the same after its rivals' exchange states. For example, if $\omega = (2, 3, 4)$, then set $\kappa = (3, 2, 1)$ in equation (21) to obtain $u_2(4, 3, 2) = u_2(\omega_3, \omega_2, \omega_1) = u_3(\omega_1, \omega_2, \omega_3) = u_2(2, 3, 4)$.

Inspection shows that these properties imply the notion of symmetry in Ericson and Pakes (1995): a firm's policy is a common function of its own state and the distribution of its rivals' states.

One of the reasons symmetry is important is that it eases the computational burden considerably. Instead of having to compute value and policy functions for *all* firms, under symmetry it suffices to compute value and policy functions for *one* firm, say firm 1. To see this, let $\kappa = (n, 2, \ldots, n-1, 1, n+1, \ldots, N)$ in equation (21) to obtain

$$u_n(\omega_n, \omega_2, \ldots, \omega_{n-1}, \omega_1, \omega_{n+1}, \ldots, \omega_N) = u_1(\omega_1, \omega_2, \ldots, \omega_{n-1}, \omega_n, \omega_{n+1}, \ldots, \omega_N), \qquad (22)$$

and similarly for the value function. That is, the value and policy of firm n is the same as the value and policy of firm 1 had their states been interchanged. In addition, symmetry reduces the size of the state space on which the value and policy functions of firm 1 are defined because firm 1 does not care about the identity of its competitors. To see this, let $n = 1$ and $\kappa = (1, 2, \ldots, k-1, l, k+1, \ldots, l-1, k, l+1, \ldots, N)$ with $k \geq 2$ and $l \geq 2$ in equation (21) to obtain

$$u_1(\omega_1, \omega_2, \ldots, \omega_l, \ldots, \omega_k, \ldots, \omega_N) = u_1(\omega_1, \omega_2, \ldots, \omega_k, \ldots, \omega_l, \ldots, \omega_N), \qquad (23)$$

and similarly for the value function. That is, only the firm's own state and the distribution of rivals' states matter. This latter property is commonly referred to as anonymity or exchangeability.[24]

We are now ready to state our symmetry assumption.

Assumption 4. The local income functions are symmetric, that is,

$$h_n(\omega_{\kappa_1}, \ldots, \omega_{\kappa_N}, u_{\kappa_1}(\omega), \ldots, u_{\kappa_N}(\omega), V_n) = h_{\kappa_n}(\omega_1, \ldots, \omega_N, u_1(\omega), \ldots, u_N(\omega), V_{\kappa_n}) \qquad (24)$$

for all $u(\omega)$, symmetric V, ω, n, and all κ.

Note that the value functions that enter the local income functions are themselves symmetric.

Some further explanation may be helpful. A permutation κ shuffles firms' states, actions, and value functions in a way that preserves the values of their local income functions according

[24] Equations (22) and (23) are often together taken as the definition of symmetry (e.g., Doraszelski and Pakes, 2007). It is easy to see that they are equivalent to our notion of symmetry in equation (21). Working with equation (21) instead of equations (22) and (23) simplifies the notation in the remainder of this section.

to the principle that identical actions in identical situations yield identical payoffs. Let $n = 2$ and $\kappa = (2, 3, 1)$ in equation (24) to obtain

$$h_2(\omega_2, \omega_3, \omega_1, u_2(\omega), u_3(\omega), u_1(\omega), V_2) = h_3(\omega_1, \omega_2, \omega_3, u_1(\omega), u_2(\omega), u_3(\omega), V_3).$$

On the left-hand side, firm 2 is in state ω_3 and takes action $u_3(\omega)$ while it faces two rivals, one in state ω_1 and one in state ω_2. On the right-hand side, firm 3 is in state ω_3 and takes action $u_3(\omega)$ while it faces two rivals, one in state ω_1 and one in state ω_2. Because the state of firm 2 on the left-hand side is that of firm 3 on the right-hand side and the distribution over states and actions of firm 2's rivals on the left-hand side is that of firm 3's rivals on the right-hand side, their respective situations are identical.

Although we have stated Assumption 4 in terms of the local income functions to facilitate the adaptation of our existence proof to other models, it is readily tied to the model's primitives.

Condition 2. The model's primitives are symmetric if (i) the profit functions are symmetric, that is,

$$\pi_n(\omega_{\kappa_1}, \ldots, \omega_{\kappa_N}) = \pi_{\kappa_n}(\omega_1, \ldots, \omega_N)$$

for all ω, n, and all κ and (ii) the transition function is symmetric, that is,

$$P(\omega'_{\kappa_1}, \ldots, \omega'_{\kappa_N}, \omega_{\kappa_1}, \ldots, \omega_{\kappa_N}, \chi_{\kappa_1}(\omega, \phi_{\kappa_1}), \ldots, \chi_{\kappa_N}(\omega, \phi_{\kappa_N}), x_{\kappa_1}(\omega), \ldots, x_{\kappa_N}(\omega))$$
$$= P(\omega'_1, \ldots, \omega'_N, \omega_1, \ldots, \omega_N, \chi_1(\omega, \phi_1), \ldots, \chi_N(\omega, \phi_N), x_1(\omega), \ldots, x_N(\omega))$$

for all $\chi(\omega, \phi), x(\omega), \omega', \omega$, *and all* κ.

Proposition 4. If the model's primitives are symmetric (Condition 2), then Assumption 4 holds.

The proof of Proposition 4 is straightforward but tedious and therefore omitted. Note that in the special case of independent transitions, part (ii) of Condition 2 is satisfied whenever the factors $P_n(\cdot)$ of the transition function $P(\cdot)$ are the same across firms, that is, $P_n(\omega'_n, \omega_n, \chi_n(\omega, \phi_n), x_n(\omega)) = P_1(\omega'_n, \omega_n, \chi_n(\omega, \phi_n), x_n(\omega))$ for all n.

Together with Assumptions 1, 2, and 3 in Section 4, Assumption 4 ensures existence of a symmetric equilibrium.

Proposition 5. Under Assumptions 1, 2, 3, and 4, a symmetric equilibrium exists in cutoff entry/exit and pure investment strategies.

The idea of the proof is as follows. Symmetry allows us to restrict attention to the best-reply correspondence of firm 1. To enforce the anonymity that symmetry implies, we redefine the state space employing Ericson and Pakes' (1995) notion that symmetry means each firm's investment is a common function of its own state and the distribution of its rivals' states. This reduced state space makes it impossible for firm 1 to tailor its policy to the identity of its competitors. An argument analogous to the proof of Proposition 2 shows that there exists a fixed point to the best-reply correspondence of firm 1. We use this fixed point to construct a candidate equilibrium by specifying symmetric policies for all firms. The associated value functions are also symmetric. Finally, to complete the argument, we exploit the symmetry of the local income functions to show that no firm has an incentive to deviate from the candidate equilibrium.

In preparation for proving Proposition 5, we introduce the necessary notation to construct the candidate equilibrium. To understand our notation, it is helpful to keep in mind that the candidate equilibrium will be symmetric. We begin with defining the reduced state space. Consider firm n and state ω. Define $\sigma_n = (\sigma_{n,1}, \ldots, \sigma_{n,M}, \sigma_{n,M+1})$, where $\sigma_{n,m}$ denotes the number of competitors of firm n that are in state m (excluding firm n), and $\Sigma = \left\{\sigma_n \in \{0, 1, \ldots, N-1\} | \sum_{m=1}^{M+1} \sigma_{n,m} = N-1\right\}$ to be the set of values that σ_n can take on. Rewrite ω as (ω_n, σ_n). Let $S^\circ = \Omega \times \Sigma$ denote the *reduced* state space and $S = \Omega^N$ the *full* state space. Define a function $\tau_n : S \rightarrow S^\circ$ such that $\tau_n(\omega) = (\omega_n, \sigma_n)$; it maps the full to the reduced state space. For example, if $N = 4$, $M = 3$, and $\omega = (3, 2, 2, 4)$, then $(\omega_1, \sigma_1) = \tau_1(\omega) = (3, 0, 2, 0, 1)$ and $(\omega_3, \sigma_3) = \tau_3(\omega) = (2, 0, 1, 1, 1)$.

Note that no information is lost in going from the full to the reduced state space, provided that the equilibrium is symmetric. In particular, $\tau_1(\omega)$ contains all the information in ω that is required to evaluate the value and policy functions of firm 1. Note also that in general the reduced state space is considerably smaller than the full state space: it has just $|S^\circ| = (M+1)\binom{M+N-1}{N-1} < (M+1)^N = |S|$ states.[25]

Define the inverse function $\tau_n^{-1} : S^\circ \to S$ such that $\omega = \tau_n^{-1}(\omega_n, \sigma_n)$ is a fixed selection from the set $\{\omega | (\omega_n, \sigma_n) = \tau_n(\omega)\}$. We adopt the convention that $\omega = \tau_n^{-1}(\omega_n, \sigma_n)$ satisfies $\omega_1 \leq \omega_2 \leq \cdots \leq \omega_{n-1} \leq \omega_{n+1} \leq \cdots \leq \omega_N$. Observe that, if $\hat{\omega} = \tau_n^{-1}(\tau_n(\omega))$, then $\hat{\omega}$ is obtained from ω by rearranging the elements of ω_{-n}. For example, $(3, 2, 2, 5) = \tau_1^{-1}(\tau_1(3, 2, 5, 2))$, $(2, 2, 3, 5) = \tau_2^{-1}(\tau_2(3, 2, 5, 2))$, and so forth. A state $\check{\omega}$ is called *canonical* if and only if $\check{\omega} = \tau_1^{-1}(\check{\omega}_1, \sigma_1)$ for some $(\check{\omega}_1, \sigma_1)$. We use the symbol ˇ to distinguish canonical states in the remainder of this section.

Next we redefine actions, strategies, and payoffs on the reduced state space. We use the symbol $\circ$ to distinguish objects defined on the reduced state space from the corresponding objects defined on the full state space. For example, we write $u_1^\circ(\omega_1, \sigma_1) \in \mathcal{U}_1^\circ(\omega_1, \sigma_1)$ instead of $u_1(\omega) \in \mathcal{U}_1(\omega)$, where $\mathcal{U}_1^\circ(\omega_1, \sigma_1) = \mathcal{U}_1(\tau_1^{-1}(\omega_1, \sigma_1))$ because $\mathcal{U}_1(\omega)$ merely hinges on ω_1 (see equation (8)). By construction, a strategy $u_1^\circ = \times_{(\omega_1,\sigma_1)\in S^\circ} u_1^\circ(\omega_1, \sigma_1) \in \times_{(\omega_1,\sigma_1)\in S^\circ} \mathcal{U}_1^\circ(\omega_1, \sigma_1) = \mathcal{U}_1^\circ$ defined on the reduced state space satisfies anonymity. Consequently, in terms of the reduced state space, a symmetric equilibrium is one in which all firms use the same strategy, that is, $u_n^\circ(\omega_n, \sigma_n) = u_1^\circ(\omega_n, \sigma_n)$ for all ω_n and all σ_n. Turning to payoffs, we take the local income function of firm 1 on the reduced state space to be

$$\begin{aligned} & h_1^\circ\big((\omega_1, \sigma_1), u_1^\circ(\omega_1, \sigma_1), u_2^\circ\big(\tau_2\big(\tau_1^{-1}(\omega_1, \sigma_1)\big)\big), \ldots, u_N^\circ\big(\tau_N\big(\tau_1^{-1}(\omega_1, \sigma_1)\big)\big), V_1^\circ\big) \\ &\quad = h_1\big(\tau_1^{-1}(\omega_1, \sigma_1), u_1^\circ(\omega_1, \sigma_1), u_2^\circ\big(\tau_2\big(\tau_1^{-1}(\omega_1, \sigma_1)\big)\big), \ldots, u_N^\circ\big(\tau_N\big(\tau_1^{-1}(\omega_1, \sigma_1)\big)\big), \Lambda_1\big(V_1^\circ\big)\big), \end{aligned} \tag{25}$$

where Λ_n maps firm 1's value (or policy) function, V_1°, defined on the reduced state space to firm n's value (or policy) function, V_n, defined on the full state space. That is, the mapping Λ_n is defined such that $V_n = \Lambda_n(V_1^\circ)$ if and only if

$$V_n(\omega) = V_1^\circ(\tau_n(\omega))$$

for all ω.

This notation permits us to define the best-reply correspondence for firm 1 and to construct the candidate equilibrium. Define the maximal return operator $H_{1,u_1^\circ}^{\circ *} : \mathcal{V}_1^\circ \to \mathcal{V}_1^\circ$ pointwise by

$$\begin{aligned} \big(H_{1,u_1^\circ}^{\circ *} V_1^\circ\big)(\omega_1, \sigma_1) = \sup_{\tilde{u}_1^\circ(\omega_1,\sigma_1)\in\mathcal{U}_1^\circ(\omega_1,\sigma_1)} & h_1^\circ\big((\omega_1, \sigma_1), \tilde{u}_1^\circ(\omega_1, \sigma_1), \\ & u_1^\circ\big(\tau_2\big(\tau_1^{-1}(\omega_1, \sigma_1)\big)\big), \ldots, u_1^\circ\big(\tau_N\big(\tau_1^{-1}(\omega_1, \sigma_1)\big)\big), V_1^\circ\big), \end{aligned}$$

where, to enforce symmetry, we take all rivals of firm 1 to use the same strategy, namely u_1°. The maximal return function $V_{1,u_1^\circ}^{\circ *}$ satisfies $V_{1,u_1^\circ}^{\circ *} = H_{1,u_1^\circ}^{\circ *} V_{1,u_1^\circ}^{\circ *}$. It is well defined and continuous in u_1° as in the proof of Proposition 2. Note that there is no circularity involved in the construction of $V_{1,u_1^\circ}^{\circ *}$ because u_1° is taken as given. Define the best-reply correspondence $\Upsilon_1^\circ : \mathcal{U}_1^\circ \to \mathcal{U}_1^\circ$ by

$$\begin{aligned} \Upsilon_1^\circ\big(u_1^\circ\big) = \Big\{ \tilde{u}_1^\circ \in \mathcal{U}_1^\circ : \tilde{u}_1^\circ(\omega_1, \sigma_1) \in \arg \sup_{\tilde{u}_1^\circ(\omega_1,\sigma_1)\in\mathcal{U}_1^\circ(\omega_1,\sigma_1)} & h_1^\circ\big((\omega_1, \sigma_1), \tilde{u}_1^\circ(\omega_1, \sigma_1), \\ u_1^\circ\big(\tau_2\big(\tau_1^{-1}(\omega_1, \sigma_1)\big)\big), \ldots, u_1^\circ\big(\tau_N\big(\tau_1^{-1}(\omega_1, \sigma_1)\big)\big), V_{1,u_1^\circ}^{\circ *}\big) & \text{ for all } (\omega_1, \sigma_1) \Big\}. \end{aligned} \tag{26}$$

Under Assumptions 1, 2, and 3, a $u_1^\circ \in \mathcal{U}_1^\circ$ exists such that $u_1^\circ \in \Upsilon_1^\circ(u_1^\circ)$. To see this note that, as in the proof of Proposition 2, $\Upsilon_1^\circ(\cdot)$ is nonempty, single valued, and upper hemicontinuous and thus a function to which Brouwer's fixed-point theorem applies.

[25] Gowrisankaran (1999b) develops an algorithm for the efficient representation of the reduced state space.

Construct a candidate equilibrium by using u_1° to define firm n's policy function on the full state space to be

$$u_n = \Lambda_n(u_1^\circ). \tag{27}$$

Turning from the equilibrium policy functions to the corresponding value functions, similarly define firm n's value function on the full state space to be

$$V^*_{n,u_{-n}} = \Lambda_n\left(V^{\circ *}_{1,u_1^\circ}\right). \tag{28}$$

By construction, the above value and policy functions are symmetric.

It remains to show that no firm has an incentive to deviate from the candidate equilibrium that, by construction, is symmetric. Specifically, we show that even if we allowed a firm to tailor its policy to the identity of its competitors (as it is always free to do in the original state space and perhaps also in reality), the firm has no incentive to do so.[26] This justifies the common practice of computing equilibria directly on the reduced state space.

Proof of Proposition 5. The proof has three steps. The first step is to show that the problem of firm n in state ω is identical to the problem of firm 1 in state $\hat{\omega}$ that is obtained by switching the first with the nth element of ω. Equation (8) implies $\mathcal{U}_n(\omega) = \mathcal{U}_1(\hat{\omega})$ so that the set of feasible actions of firm n in state ω is the same as that of firm 1 in state $\hat{\omega}$. Moreover, for an arbitrary action $\tilde{u}_n(\omega) \in \mathcal{U}_n(\omega)$, we have

$$\begin{aligned} &h_n\left(\omega, u_1(\omega), u_2(\omega), \ldots, u_{n-1}(\omega), \tilde{u}_n(\omega), u_{n+1}(\omega), \ldots, u_N(\omega), V^*_{n,u_{-n}}\right) \\ &\quad = h_1\left(\hat{\omega}, \tilde{u}_n(\omega), u_2(\omega), \ldots, u_{n-1}(\omega), u_1(\omega), u_{n+1}(\omega), \ldots, u_N(\omega), V^*_{1,u_{-1}}\right) \\ &\quad = h_1\left(\hat{\omega}, \tilde{u}_n(\omega), u_2(\hat{\omega}), \ldots, u_{n-1}(\hat{\omega}), u_n(\hat{\omega}), u_{n+1}(\hat{\omega}), \ldots, u_N(\hat{\omega}), V^*_{1,u_{-1}}\right), \end{aligned}$$

where the first equality follows from the symmetry of the value and local income functions and the second from the symmetry of the policy functions. Hence, the local income function of firm n in state ω is the same as that of firm 1 in state $\hat{\omega}$.

The second step is to show that the problem of firm 1 in the (possibly) noncanonical state $\hat{\omega}$ is identical to the problem of firm 1 in the canonical state $\check{\omega}$ that is obtained from $\hat{\omega}$ by rearranging the elements of $\hat{\omega}_{-1}$. Formally, $\check{\omega}_1 = \hat{\omega}_1$ and $\check{\omega}_n = \hat{\omega}_{\kappa_n}$ for some permutation $\kappa_{-1} = (\kappa_2, \ldots, \kappa_N)$ of $(2, \ldots, N)$. We have $\mathcal{U}_1(\hat{\omega}) = \mathcal{U}_1(\check{\omega})$ for the set of feasible actions and, for an arbitrary action $\tilde{u}_1(\hat{\omega}) \in \mathcal{U}_1(\hat{\omega})$,

$$\begin{aligned} &h_1\left(\hat{\omega}, \tilde{u}_1(\hat{\omega}), u_2(\hat{\omega}), \ldots, u_N(\hat{\omega}), V^*_{1,u_{-1}}\right) \\ &\quad = h_1\left(\check{\omega}, \tilde{u}_1(\hat{\omega}), u_{\kappa_2}(\hat{\omega}), \ldots, u_{\kappa_N}(\hat{\omega}), V^*_{1,u_{-1}}\right) \\ &\quad = h_1\left(\check{\omega}, \tilde{u}_1(\hat{\omega}), u_2(\check{\omega}), \ldots, u_N(\check{\omega}), V^*_{1,u_{-1}}\right), \end{aligned}$$

where the first equality follows from the symmetry of the value and local income functions and the second from the symmetry of the policy functions.

The third and final step is to show that firm 1 in the canonical state $\check{\omega}$ has no incentive to deviate from the candidate equilibrium. For an arbitrary action $\tilde{u}_1(\check{\omega}) \in \mathcal{U}_1(\check{\omega})$, we have

$$\begin{aligned} &h_1\left(\check{\omega}, \tilde{u}_1(\check{\omega}), u_2(\check{\omega}), \ldots, u_N(\check{\omega}), V^*_{1,u_{-1}}\right) \\ &\quad = h_1\left(\check{\omega}, \tilde{u}_1(\check{\omega}), u_1^\circ(\tau_2(\check{\omega})), \ldots, u_1^\circ(\tau_N(\check{\omega})), \Lambda_1\left(V^{\circ *}_{1,u_1^\circ}\right)\right) \\ &\quad = h_1\left(\tau_1^{-1}(\check{\omega}_1, \sigma_1), \tilde{u}_1(\check{\omega}), u_1^\circ\left(\tau_2\left(\tau_1^{-1}(\check{\omega}_1, \sigma_1)\right)\right), \ldots, u_1^\circ\left(\tau_N\left(\tau_1^{-1}(\check{\omega}_1, \sigma_1)\right)\right), \Lambda_1\left(V^{\circ *}_{1,u_1^\circ}\right)\right) \\ &\quad = h_1^\circ\left((\check{\omega}_1, \sigma_1), \tilde{u}_1(\check{\omega}), u_1^\circ\left(\tau_2\left(\tau_1^{-1}(\check{\omega}_1, \sigma_1)\right)\right), \ldots, u_1^\circ\left(\tau_N\left(\tau_1^{-1}(\check{\omega}_1, \sigma_1)\right)\right), V^{\circ *}_{1,u_1^\circ}\right), \end{aligned} \tag{29}$$

[26] This final step is absent from Pesendorfer and Schmidt-Dengler's (2008) proof of their Corollary 1 asserting that a symmetric equilibrium exists if the primitives are symmetric.

where the first equality follows from equations (27) and (28), the second from the fact that $\breve{\omega} = \tau^{-1}(\breve{\omega}_1, \sigma_1)$ for some $(\breve{\omega}_1, \sigma_1)$ because $\breve{\omega}$ is canonical, and the last from equation (29). Moreover, we have $\mathcal{U}_1(\breve{\omega}) = \mathcal{U}_1(\tau_1^{-1}(\breve{\omega}_1, \sigma_1)) = \mathcal{U}_1^{\circ}(\breve{\omega}_1, \sigma_1)$ for the set of feasible actions. Because the last line of equation (29) is the maximand of firm 1 in the best-reply correspondence in equation (26), firm 1 has no incentive to deviate. Moreover, because the problem of firm n in state ω is identical to the problem of firm 1 in state $\breve{\omega}$ by the first two steps of the proof, no firm has an incentive to deviate from the candidate equilibrium. *Q.E.D.*

Combining Propositions 3, 4, and 5, we are ready to state our main result establishing that a computationally tractable equilibrium exists in our model.

Theorem 1. Suppose Assumptions 1 and 2 hold. If the transition function $P(\cdot)$ is UIC admissible (Condition 1) and the model's primitives are symmetric (Condition 2), then a symmetric equilibrium exists in cutoff entry/exit and pure investment strategies.

7. Convergence to equilibria in mixed strategies

■ In this section, we relate our game with random scrap values/setup costs to the game of complete information. To do so, we write firm n's scrap value as $\phi + \epsilon\theta_n$ if $\omega_n \neq M+1$ and its setup cost as $\phi^e + \epsilon\theta_n^e$ if $\omega_n = M+1$, where $\epsilon > 0$ is a constant scale factor that measures the importance of incomplete information. Overloading notation, we assume that $\theta_n \sim F(\cdot)$ and $\theta_n^e \sim F^e(\cdot)$ with $\mathrm{E}(\theta_n) = \mathrm{E}(\theta_n^e) = 0$. Substituting into equation (10), firm n's return or local income function $h_n^{\epsilon}(\cdot)$ becomes

$$h_n^{\epsilon}(\omega, u_n(\omega), V_n) = \begin{cases} \pi_n(\omega) + (1-\xi_n(\omega))\phi + \epsilon \int_{\theta_n > F^{-1}(\xi_n(\omega))} \theta_n dF(\theta_n) & \\ \quad + \xi_n(\omega)\Big\{-x_n(\omega) + \beta \mathrm{E}\{V_n(\omega')|\omega, \omega_n' \neq M+1, \xi_{-n}(\omega), x(\omega)\}\Big\} & \text{if } \omega_n \neq M+1, \\ -\epsilon \int_{\theta_n^e < F^{e-1}(\xi_n(\omega))} \theta_n^e dF^e(\theta_n^e) & \\ \quad + \xi_n(\omega)\Big\{-\phi^e - x_n(\omega) + \beta \mathrm{E}\{V_n(\omega')|\omega, \omega_n' \neq M+1, \xi_{-n}(\omega), x(\omega)\}\Big\} & \text{if } \omega_n = M+1, \end{cases}$$

where $\xi_n(\omega) = \int \chi_n(\omega, \theta_n) dF(\theta_n) = \int 1(\phi + \epsilon\theta_n < \bar{\phi}_n(\omega)) dF(\theta_n) = F(\frac{\bar{\phi}_n(\omega) - \phi}{\epsilon})$, and so forth.

Proposition 2 in Section 4 guarantees the existence of an equilibrium in cutoff entry/exit and pure investment strategies for any fixed $\epsilon > 0$. Note that $h_n^0(\cdot)$ is the local income function that obtains in a game of complete information. As our example in Section 3 has shown, there is a need to allow for mixed entry/exit strategies in a game with deterministic scrap values/setup costs such as in Ericson and Pakes (1995). We thus ask whether the equilibrium of the game of incomplete information converges to the equilibrium in mixed entry/exit strategies as ϵ approaches zero. The following proposition gives an affirmative answer.

Proposition 6. Suppose Assumptions 1, 2, and 3 hold, and consider a sequence $\{\epsilon^l\}$ such that $\lim_{l\to\infty} \epsilon^l = 0$. Let $\{u^l\}$ be a corresponding sequence of equilibria in cutoff entry/exit strategies such that $\lim_{l\to\infty} u^l = u$. Then u is an equilibrium in mixed entry/exit strategies.

Proof. Let $\{V_{u^l}^{\epsilon^l}\}$ be the corresponding sequence of return functions where $V_{n,u^l}^{\epsilon^l}$ satisfies $V_{n,u^l}^{\epsilon^l} = H_{n,u^l}^{\epsilon^l} V_{n,u^l}^{\epsilon^l}$. Repeating the argument that led to equation (15) in Section 4 shows that each element of $\{V_{u^l}^{\epsilon^l}\}$ is well defined due to Assumption 1. Moreover, because $H_{n,u}^{\epsilon} V_n$ is continuous in ϵ and u for all V_n, Whitt (1980) implies that the return function $V_{n,u}^{\epsilon}$ is continuous in ϵ and u. Let $V_{n,u} = \lim_{l\to\infty} V_{n,u^l}^{\epsilon^l}$ for all n.

The proof proceeds in two steps. In the first step, we verify that the limiting strategy u_n is optimal given the return function $V_{n,u}$ for all n. In the second step, we verify that the return function $V_{n,u}$ coincides with the maximal return function for all n.

Suppose $u_n(\omega) \notin \arg\sup_{\tilde{u}_n(\omega)\in\mathcal{U}_n(\omega)} h_n^0(\omega, \tilde{u}_n(\omega), u_{-n}(\omega), V_{n,u})$ for some ω and some n. Then there exists $\tilde{u}_n(\omega) \in \mathcal{U}_n(\omega)$ such that

$$h_n^0(\omega, \tilde{u}_n(\omega), u_{-n}(\omega), V_{n,u}) > h_n^0(\omega, u_n(\omega), u_{-n}(\omega), V_{n,u}).$$

Because $h_n^{\epsilon}(\omega, u(\omega), V_{n,u})$ is a continuous function of ϵ, $u(\omega)$, and $V_{n,u}$, there exists L large enough such that

$$h_n^{\epsilon^l}\left(\omega, \tilde{u}_n(\omega), u_{-n}^l(\omega), V_{n,u^l}^{\epsilon^l}\right) > h_n^{\epsilon^l}\left(\omega, u_n^l(\omega), u_{-n}^l(\omega), V_{n,u^l}^{\epsilon^l}\right)$$

for all $l \geq L$. Hence, $u_n^l(\omega) \notin \arg\sup_{\tilde{u}_n(\omega)\in\mathcal{U}_n(\omega)} h_n^{\epsilon^l}(\omega, \tilde{u}_n(\omega), u_{-n}^l(\omega), V_{n,u^l}^{\epsilon^l})$ and we obtain a contradiction.

It remains to verify that the return function $V_{n,u}$ coincides with the maximal return function for all n. By construction, $V_{n,u^l}^{\epsilon^l}$ satisfies $V_{n,u^l}^{\epsilon^l}(\omega) = h_n^{\epsilon^l}(\omega, u^l(\omega), V_{n,u^l}^{\epsilon^l})$ for all ω. Taking limits on both sides shows that $V_{n,u}$ satisfies $V_{n,u}(\omega) = h_n^0(\omega, u(\omega), V_{n,u})$ for all ω. Using the first step of the proof, we have

$$V_{n,u}(\omega) = h_n^0(\omega, u(\omega), V_{n,u}) = \sup_{\tilde{u}_n(\omega)\in\mathcal{U}_n(\omega)} h_n^0(\omega, \tilde{u}_n(\omega), u_{-n}(\omega), V_{n,u})$$

for all ω. Because $V_{n,u}$ is a fixed point of the maximal return operator of the game of complete information, it is the maximal return function. *Q.E.D.*

Convergence results for static games date back at least to Harsanyi (1973) but, to the best of our knowledge, ours is the first such result for dynamic stochastic games.[27] The proof of Proposition 6 relies on the continuity of the return function $V_{n,u}^{\epsilon}$. The fact that continuity obtains further illustrates the power of the dynamic programming approach.

Note that Proposition 6 does not imply that $\lim_{l\to\infty} u^l$ exists. On the other hand, because $\mathcal{U}$ is compact, every sequence $\{u^l\}$ has a convergent subsequence, and Proposition 6 applies to the subsequential limit. This establishes the following.

Corollary 1. Under Assumptions 1, 2, and 3, an equilibrium exists in mixed entry/exit and pure investment strategies in the Ericson and Pakes' (1995) model.

8. Conclusions

■ This article provides a general model of dynamic competition in an oligopolistic industry with investment, entry, and exit and ensures that there exists a computationally tractable equilibrium for it. Our starting point is the observation that existence of an equilibrium in the Ericson and Pakes' (1995) game of complete information requires mixed entry/exit strategies. This is problematic from a computational point of view because the existing algorithms—notably Pakes and McGuire (1994, 2001)—cannot cope with mixed strategies. We therefore introduce firm heterogeneity in the form of randomly drawn, privately known scrap values and setup costs into the model. We show that the resulting game of incomplete information always has an equilibrium in cutoff entry/exit strategies that is no more demanding to compute than a (possibly nonexistent) equilibrium in pure entry/exit strategies of the original game of complete information. We further ensure that the equilibrium is in pure investment strategies by first assuming that a firm's investment choice always is uniquely determined. We then show that this assumption is satisfied provided the transition function is UIC admissible. This, in fact, is a key contribution because UIC admissibility is defined with respect to the model's primitives and is easily checked.

We build on our basic existence result in three ways. First, we show that a symmetric equilibrium exists under the appropriate assumptions on the model's primitives. Requiring the

[27] In subsequent work, Doraszelski and Escobar (2010) provide a convergence result for general dynamic stochastic games with finite state and action spaces. They also show that the approachability part of Harsanyi's (1973) purification theorem carries over from static games to dynamic stochastic games. That is, all equilibria of the original game are approached by some equilibrium of the perturbed game as the perturbation vanishes.

equilibrium to be symmetric is important because it reduces the computational burden and forces heterogeneity to arise endogenously among *ex ante* identical firms. Second, we show that, as the distribution of the random scrap values/setup costs becomes degenerate, equilibria in cutoff entry/exit strategies converge to equilibria in mixed entry/exit strategies of the game of complete information. Third, as a byproduct, this last result implies that there exists an equilibrium in the Ericson and Pakes' (1995) model, provided that mixed entry/exit strategies are admissible.

References

AGUIRREGABIRIA, V. AND MIRA, P. "Sequential Estimation of Dynamic Discrete Games." Working Paper, CEMFI, 2004.

—— AND ——. "Sequential Estimation of Dynamic Discrete Games." *Econometrica*, Vol. 75 (2007), pp. 1–54.

AMIR, R. "Continuous Stochastic Games of Capital Accumulation with Convex Transitions." *Games and Economic Behavior*, Vol. 15 (1996), pp. 111–131.

BAJARI, P., BENKARD, L., AND LEVIN, J. "Estimating Dynamic Models of Imperfect Competition." *Econometrica*, Vol. 75 (2007), pp. 1331–1370.

BENKARD, L. "A Dynamic Analysis of the Market for Wide-Bodied Commercial Aircraft." *Review of Economic Studies*, Vol. 71 (2004), pp. 581–611.

BERESTEANU, A. AND ELLICKSON, P. "The Dynamics of Retail Oligopoly." Working Paper, Duke University, 2006.

BESANKO, D. AND DORASZELSKI, U. "Capacity Dynamics and Endogenous Asymmetries in Firm Size." *RAND Journal of Economics*, Vol. 35 (2004), pp. 23–49.

——, ——, KRYUKOV, Y., AND SATTERTHWAITE, M. "Learning-by-Doing, Organizational Forgetting, and Industry Dynamics." *Econometrica*, Vol. 78 (2010), pp. 453–508.

——, ——, LU, L., AND SATTERTHWAITE, M. "Lumpy Capacity Investment and Disinvestment Dynamics." *Operations Research* (2010), forthcoming.

BLACKWELL, D. "Discounted Dynamic Programming." *Annals of Mathematical Statistics*, Vol. 36 (1965), pp. 226–235.

CHAKRABARTI, S. "Markov Equilibria in Discounted Stochastic Games." *Journal of Economic Theory*, Vol. 85 (1999), pp. 392–327.

——. "Pure Strategy Markov Equilibrium in Stochastic Games with a Continuum of Players." *Journal of Mathematical Economics*, Vol. 39 (2003), pp. 693–724.

COLLARD-WEXLER, A. "Demand Fluctuations and Plant Turnover in Ready-Mix Concrete." Working Paper, New York University, 2006.

CURTAT, L. "Markov Equilibria of Stochastic Games with Complementarities." *Games and Economic Behavior*, Vol. 17 (1996), pp. 177–199.

DENARDO, E. "Contraction Mappings in the Theory Underlying Dynamic Programming." *SIAM Review*, Vol. 9 (1967), pp. 165–177.

DORASZELSKI, U. AND ESCOBAR, J. "A Theory of Regular Markov Perfect Equilibria in Dynamic Stochastic Games: Genericity, Stability, and Purification." *Theoretical Economics* (2010), forthcoming.

—— AND MARKOVICH, S. "Advertising Dynamics and Competitive Advantage." *RAND Journal of Economics*, Vol. 38 (2007), pp. 557–592.

—— AND PAKES, A. "A Framework for Applied Dynamic Analysis in IO." In M. Armstrong and R. Porter, eds., *Handbook of Industrial Organization*, Vol. 3. Amsterdam: North-Holland, 2007.

—— AND SATTERTHWAITE, M. "Foundations of Markov-Perfect Industry Dynamics: Existence, Purification, and Multiplicity." Working Paper no. 1383, Center for Mathematical Studies in Economics and Management Science, Northwestern University, 2003.

DUTTA, P. AND SUNDARAM, R. "Markovian Equilibrium in a Class of Stochastic Games: Existence Theorems for Discounted and Undiscounted Models." *Economic Theory*, Vol. 2 (1992), pp. 197–214.

ERICSON, R. AND PAKES, A. "Markov-Perfect Industry Dynamics: A Framework for Empirical Work." *Review of Economic Studies*, Vol. 62 (1995), pp. 53–82.

ESCOBAR, J. "Existence of Pure and Behavior Strategy Stationary Markov Equilibrium in Dynamic Stochastic Games." Working Paper, Stanford University, 2008.

FEDERGRUEN, A. "On *N*-Person Stochastic Games with Denumerable State Space." *Advances in Applied Probability*, Vol. 10 (1978), pp. 452–471.

FINK, A. "Equilibrium Points of Stochastic Noncooperative Games." *Journal of Science of the Hiroshima University Series A*, Vol. 28 (1964), pp. 89–93.

FUDENBERG, D., GILBERT, R., STIGLITZ, J., AND TIROLE, J. "Preemption, Leapfrogging and Competition in Patent Races." *European Economic Review*, Vol. 22 (1983), pp. 3–31.

GOWRISANKARAN, G. "A Dynamic Analysis of Mergers." Ph.D. Thesis, Department of Economics, Yale University, 1995.

——. "A Dynamic Model of Endogenous Horizontal Mergers." *RAND Journal of Economics*, Vol. 30 (1999a), pp. 56–83.

——. "Efficient Representation of State Spaces for Some Dynamic Models." *Journal of Economic Dynamics and Control*, Vol. 23 (1999b), pp. 1077–1098.

—— AND TOWN, R. "Dynamic Equilibrium in the Hospital Industry." *Journal of Economics and Management Strategy*, Vol. 6 (1997), pp. 45–74.

HARRIS, C. AND VICKERS, J. "Racing with Uncertainty." *Review of Economic Studies*, Vol. 54 (1987), pp. 1–21.

HARSANYI, J. "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points." *International Journal of Game Theory*, Vol. 2 (1973), pp. 1–23.

HORST, U. "Stationary Equilibria in Discounted Stochastic Games with Weakly Interacting Players." *Games and Economic Behavior*, Vol. 51 (2005), pp. 83–108.

JENKINS, M., LIU, P., MATZKIN, R., AND MCFADDEN, D. "The Browser War: Econometric Analysis of Markov Perfect Equilibrium in Markets with Network Effects." Working Paper, Stanford University, 2004.

JOFRE-BONET, M. AND PESENDORFER, M. "Estimation of a Dynamic Auction Game." *Econometrica*, Vol. 71 (2003), pp. 1143–1489.

MASKIN, E. AND TIROLE, J. "A Theory of Dynamic Oligopoly, III: Cournot Competition." *European Economic Review*, Vol. 31 (1987), pp. 947–968.

—— AND ——. "A Theory of Dynamic Oligopoly, I: Overview and Quantity Competition with Large Fixed Costs." *Econometrica*, Vol. 56 (1988a), pp. 549–569.

—— AND ——. "A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles." *Econometrica*, Vol. 56 (1988b), pp. 571–599.

—— AND ——. "Markov Perfect Equilibrium, I: Observable Actions." *Journal of Economic Theory*, Vol. 100 (2001), pp. 191–219.

MCKELVEY, R. AND MCLENNAN, A. "Computation of Equilibria in Finite Games." In H. AMMAN, D. KENDRICK, AND J. RUST, eds., *Handbook of Computational Economics*. Amsterdam: North-Holland, 1996.

MERTENS, J. "Stochastic Games." In R. Aumann and S. Hart, eds., *Handbook of Game Theory*, Vol. 3. Amsterdam: Elsevier, 2002.

NOWAK, A. "On Stochastic Games in Economics." *Mathematical Methods of Operations Research*, Vol. 66 (2007), pp. 513–530.

PAKES, A. AND MCGUIRE, P. "Computing Markov-Perfect Nash Equilibria: Numerical Implications of a Dynamic Differentiated Product Model." *RAND Journal of Economics*, Vol. 25 (1994), pp. 555–589.

—— AND ——. "Stochastic Algorithms, Symmetric Markov Perfect Equilibrium, and the 'Curse' of Dimensionality." *Econometrica*, Vol. 69 (2001), pp. 1261–1281.

——, OSTROVSKY, M., AND BERRY, S. "Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry/Exit Examples)." *RAND Journal of Economics*, Vol. 38 (2007), pp. 373–399.

PESENDORFER, M. AND SCHMIDT-DENGLER, P. "Identification and Estimation of Dynamic Games." Working Paper no. 9726, NBER, 2003.

—— AND ——. "Asymptotic Least Squares Estimators for Dynamic Games." *Review of Economic Studies*, Vol. 75 (2008), pp. 901–928.

RYAN, S. "The Costs of Environmental Regulation in a Concentrated Industry." Working Paper, MIT, 2006.

SOBEL, M. "Noncooperative Stochastic Games." *Annals of Mathematical Statistics*, Vol. 42 (1971), pp. 1930–1935.

STOKEY, N. AND LUCAS, R., WITH PRESCOTT, E. *Recursive Methods in Economic Dynamics*. Cambridge, Mass.: Harvard University Press, 1989.

WHITT, W. "Representation and Approximation of Noncooperative Sequential Games." *SIAM Journal of Control and Optimization*, Vol. 18 (1980), pp. 33–48.