

Computable Markov-Perfect Industry Dynamics: Existence, Purification, and Multiplicity*

Ulrich Doraszelski

Department of Economics, Harvard University, and CEPR[†]

Mark Satterthwaite

Kellogg School of Management, Northwestern University[‡]

September 26, 2007

Abstract

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 & Pakes (1995) model of industry dynamics requires admissibility of mixed entry/exit strategies, contrary to the assertion in their paper, 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. Finally, we provide the first example of multiple symmetric equilibria in this literature.

1 Introduction

Over the past few years, the industrial organization literature has made considerable progress in analyzing the dynamics of industry. In a seminal paper, Ericson & Pakes (1995) provide a 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

*We are indebted to Lanier Benkard, David Besanko, Michaela Draganska, Dino Gerardi, Gautam Gowrisankaran, Ken Judd, Patricia Langohr, Andy Skrzypacz, and Eilon Solan for helpful discussions. 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.

[†]Cambridge, MA 02138, U.S.A., doraszelski@harvard.edu.

[‡]Evanston, IL 60208, U.S.A., m-satterthwaite@kellogg.northwestern.edu.

date have yielded novel insights into a number of important questions in industrial organization (e.g., Berry & Pakes 1993, Gowrisankaran 1999*a*, Fershtman & Pakes 2000) and it provides a starting point for ongoing research in industrial organization (see Doraszelski & Pakes (2007) for a survey) as well as in other fields such as international trade (Erdem & Tybout 2003) and finance (Goettler, Parlour & Rajan 2005, Kadyrzhanova 2006). More recently, Aguirregabiria & Mira (2007), Bajari, Benkard & Levin (2007), Pakes, Ostrovsky & Berry (2006), and Pesendorfer & Schmidt-Dengler (2003) 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 & Town 1997, Benkard 2004, Beresteanu & Ellickson 2005, Collard-Wexler 2005, Ryan 2005).

To achieve this goal the researcher has to be able to compute the Markov perfect equilibrium using the estimated primitives. This, in turn, requires that an equilibrium exists. Unfortunately, existence can not be guaranteed under the conditions in Ericson & Pakes (1995).

Moreover, existence by itself is not enough for two reasons. First, contrary to the assertion in their paper, the existence of an equilibrium in the Ericson & 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 despite the considerable progress that has been made in the context of finite games (see McKelvey & McLennan (1996) for a survey). Indeed, the algorithms developed by Pakes & McGuire (1994, 2001) cannot cope with mixed entry/exit strategies. Worse, computing mixed strategies over continuous actions such as investment is completely infeasible at present. 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 an equilibrium in pure strategies that is symmetric.

Our goal in this paper is to modify the Ericson & Pakes (1995) model just enough to ensure that there exists a computationally tractable equilibrium for it. Thus, just as in the Ericson & Pakes (1995) model, in our model incumbent firms decide in each period whether to remain in the industry and how much to invest, and potential entrants decide whether to enter the industry. Once the investment, entry, and exit decisions are made, firms compete in the product market. All characteristics that are relevant to a firm’s profit from product market competition are encoded in its “state,” and the firm is able to change its state over time through investment. For example, a firm’s state may describe its production capacity, cost structure, or the quality of its product.

We reformulate the Ericson & Pakes (1995) model as a game of incomplete information. We show that doing so does not alter the fundamental economics of the model. It does, however, allow us to prove the existence of an equilibrium in pure strategies. We further

show that under reasonable conditions the equilibrium is symmetric. This fulfills our goal of establishing that the dynamic industry equilibrium is computationally tractable.

A further goal of this paper is to provide a step-by-step guide to formulating dynamic models of industry equilibrium that are both computationally tractable and theoretically rigorous. There are three difficulties that we now discuss in detail.

Cutoff entry/exit strategies. First, the existence of an equilibrium cannot be ensured without allowing firms to randomize, in some way or another, over discrete actions such as entry and exit. Since Ericson & 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).

The game-theoretic literature has, of course, long recognized the importance of allowing firms to randomize. Unfortunately, the existence theorems in the extant literature rely on computationally intractable mixed strategies (see Mertens (2002) for a survey). Strictly speaking, these existence theorems are not even applicable because they cover dynamic stochastic games with either discrete (e.g., Fink 1964, Sobel 1971, Maskin & Tirole 2001) or continuous actions (e.g., Federgruen 1978, Whitt 1980), whereas the Ericson & Pakes (1995) model combines discrete entry/exit decisions with continuous investment decisions. Lastly, they neglect the requirement of symmetry. In short, the existence theorems in the game-theoretic literature are ill-suited for our purposes.

To eliminate the need for mixed entry/exit strategies without jeopardizing existence, we extend Harsanyi's (1973*a*) 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, i.e., while 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 & 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 if* it was following a mixed strategy. Note that random 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*a*) 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.

Pakes & McGuire (1994) suggest treating a potential entrant’s setup cost as a random variable to overcome convergence problems in their algorithm.¹ To our knowledge Gowrisankaran (1995) is the first to make the connection between existence and randomization of both entry and exit decisions. In his dynamic endogenous merger model he assumes randomly drawn, privately known setup costs and scrap values but then goes on to write down the Bellman equation of an incumbent firm *as if* scrap values were deterministic. However, this equation, unless written down in error, fails to define a contraction and thus renders inapplicable the dynamic programming techniques that are the basis of the Ericson & Pakes (1995) model. Benkard (2004) also assumes randomly drawn, privately known setup costs and scrap values in his model of dynamic competition among makers of wide-body jet aircraft and incorporates these into his numerical algorithm in exactly the manner that we propose in this paper. Thus our existence result applies to his model. Inexplicably the Bellman equation of an incumbent firm that he writes down makes it appear as if he treats scrap values as deterministic. This, however, is a proof reading error.²

In summary, while the idea of using random setup costs/scrap values instead of mixed entry/exit strategies is part of the folklore in the literature following Ericson & Pakes (1995), we do not know of a coherent treatment of the resulting game of incomplete information. This is unfortunate because mistakes are easily made in formalizing a model of dynamic competition and have potentially severe consequences. Moreover, the question remains whether the “computational trick” of using random setup costs/scrap values changes the nature of strategic interactions among firms. We show in section 4.3 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 of the complete-information game. 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 there exists an equilibrium in the Ericson & Pakes (1995) model provided that mixed entry/exit strategies are admissible.

Pure investment strategies. The second difficulty is to ensure pure investment strategies. The extant game-theoretic literature routinely allows for randomization over continu-

¹Following Rust (1994) the empirical literature has used randomness to “smooth out” discontinuities caused by the discreteness of entry/exit decisions (e.g., Seim 2001). In independent work Pesendorfer & Schmidt-Dengler (2003) use randomly drawn, privately known shocks to establish the existence of a Markov perfect equilibrium in cutoff strategies in dynamic stochastic games with discrete actions (e.g., entry/exit games). Their proof, which is quite elegant, requires that the shocks are unbounded, so that all actions are played with positive probability at all times. In our setting this assumption seems awkward as it essentially stipulates that someone is willing to pay an unbounded amount to acquire the assets of a firm that makes bounded profits from product market competition.

²In June 2007 Lanier Benkard kindly wrote to us about our discussion of his paper in a previous version of our paper. In his communication he agreed that his equation (1) on p. 587 as written is incorrect, but asserted that he, in fact, had correctly calculated the equilibrium. To support this assertion he enclosed excerpts from his code. Inspection convinced us that he did the calculation in the manner we argue is appropriate.

ous actions. Computing mixed strategies over continuous actions, however, is not practical. We therefore must make sure that a firm’s optimal investment level is always unique, for that guarantees that the equilibrium is in pure investment strategies. 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. UIC admissibility is an easily verifiable condition on the model’s primitives and is not overly limiting. Indeed, while the transition functions used in the vast majority of applications of Ericson & Pakes’s (1995) framework are UIC admissible, they all have restricted a firm to transit to immediately adjacent states. Our condition establishes that this is unnecessarily restrictive, and we show how to specify more general UIC admissible transition functions.

We emphasize that UIC admissibility is a sufficient condition and that, 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 (2005) and Besanko, Doraszelski, Lu & Satterthwaite (2006) extend our handling of entry and exit to the case of discrete (or “lumpy”) investment. Their models remain computationally tractable because the equilibrium is in cutoff investment strategies. Focusing on the case of continuous investment, Jenkins, Liu, Matzkin & 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 (see their assumption 4 and theorem 2 in appendix 1). Again its rivals perceive the firm *as if* it was 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.

Our paper furthermore contributes to a growing literature that establishes the existence of a Markov perfect equilibrium in pure strategies for a variety of dynamic stochastic games whose structures are tailored to represent situations of economic interest. Curtat (1996) does so in a game with a continuum of states by assuming that the per-period payoffs as well as the transition distribution function satisfy monotonicity, supermodularity, and dominant-diagonal conditions. This entails restrictions on how per-period payoffs can vary with the state whereas our approach accommodates arbitrary per-period payoffs. Chakrabarti (2003) shows that there exists a Markov perfect equilibrium in pure strategies in a dynamic stochastic game with a continuum of players provided that the per-period payoffs and the transition density function depend only on the *average* response of the players. In subsequent work Escobar (2007) establishes the existence of a Markov perfect equilibrium in pure strategies in a 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. Since 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 & Pakes (1995) model. However, there is no systematic treatment of incomplete information as a means to purify the discrete entry/exit decisions.

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. In section 4.2 we show that this is the case under appropriate assumptions on the model’s primitives. 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 must arise endogenously from the idiosyncratic outcomes that the *ex ante* identical firms realize from their investments.

Resolving these difficulties allows us to establish that a computationally tractable symmetric equilibrium in pure strategies always exists. Given this, an important question is whether or not it is unique. To our knowledge, all applications of Ericson & Pakes’s (1995) framework have found a single equilibrium. In fact, it is often held that “nonuniqueness does not seem to be a problem” in this setting (Pakes & McGuire 1994, p. 570). We settle this issue by providing three examples of multiple symmetric equilibria. Our examples in turn highlight firms’ investment decisions, their entry/exit decisions, and product market competition as a possible source of multiple equilibria.

A minimalist view of our paper is that it employs standard techniques to rebuild the shaky foundations of an important model. In particular, random setup costs have already been suggested by Pakes & McGuire (1994) and both Gowrisankaran (1995) and Benkard (2004) state that they have taken up this suggestion in their studies of the hospital and airframe industries. But neither they nor any other authors wrote down precisely how to implement this reasonable suggestion and how it actually works to ensure existence. Consequently, learning how to use the Ericson & Pakes (1995) framework has been an unnecessarily risky and error-prone undertaking, and we have ourselves wasted much time when we initially incorrectly formulated the model. Our motivation for writing this paper is therefore to provide a 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 step-by-step 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 paper is as follows. We develop the model in section 2. In section 3 we provide a series of simple examples to illustrate the key themes of the subsequent analysis and discuss a number of computational issues. We turn to the analysis of the full model in section 4 where we show that a computationally tractable symmetric equilibrium in pure strategies always exists in the reformulated Ericson & Pakes (1995) model and that equilibria of the incomplete-information game converge to equilibria of the complete-information game. In section 5 we introduce UIC admissibility and explain how to specify general transition functions that satisfy this criterion. In section 6 we provide several examples of multiple equilibria. Section 7 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. While 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, 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). Since 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 while state $M+1$ identifies the firm as being inactive.³ At any point in time the industry is completely characterized by the list of firms' states $\omega = (\omega_1, \dots, \omega_N) \in S$

³This 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.

where $S = \Omega^N$ is the state space.⁴ 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.⁵ 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:

1. Incumbent firms learn their scrap value and decide on exit and investment. Potential entrants learn their setup cost and decide on entry and investment.
2. Incumbent firms compete in the product market.
3. Exit and entry decisions are implemented.
4. 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.

While 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. Hence, 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 from a distribution $F(\cdot)$ with $E(\phi_n) = \phi$.⁶ Scrap values are independently and identically

⁴Time-varying characteristics of the competitive environment are easily added to the description of the industry. Besanko & 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.

⁵Limiting 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.

⁶It is straightforward to allow firm n ’s scrap value ϕ_n to vary systematically with its state ω_n by replacing $F(\cdot)$ by $F_{\omega_n}(\cdot)$.

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 $\chi_n(\omega, \phi_n) = 1$ 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. Since 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 & 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 the 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 from a distribution $F^e(\cdot)$ with $E(\phi_n^e) = \phi^e$. Like scrap values, setup costs are independently and identically distributed across firms and periods, and its setup cost is private to a 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 & 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 & Pakes (1995) where the entrant is

randomly assigned to an arbitrary position and thus has no control over its initial position within the industry.⁷

We make these two changes because industry evolution frequently takes the form of a preemption race (Fudenberg, Gilbert, Stiglitz & Tirole 1983, Harris & Vickers 1987, Besanko & Doraszelski 2004). 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.

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 $P : S^2 \times \{0, 1\}^{2N} \times [0, \bar{x}]^N \times [0, \bar{x}^e]^N \rightarrow [0, 1]$. Thus, $P(\omega', \omega, \chi(\omega, \phi), \chi^e(\omega, \phi^e), x(\omega), x^e(\omega))$ is the probability that the industry moves from state ω to state ω' given that the incumbent firms’ exit decision are $\chi(\omega, \phi) = (\chi_1(\omega, \phi_1), \dots, \chi_N(\omega, \phi_N))$, their investment decisions are $x(\omega) = (x_1(\omega), \dots, x_N(\omega))$, etc. Necessarily $P(\omega', \omega, \chi(\omega, \phi), \chi^e(\omega, \phi^e), x(\omega), x^e(\omega)) \geq 0$ and $\sum_{\omega' \in S} P(\omega', \omega, \chi(\omega, \phi), \chi^e(\omega, \phi^e), x(\omega), x^e(\omega)) = 1$.

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

$$\prod_{\substack{n=1, \dots, N, \\ \omega_n \neq M+1}} P_n(\omega'_n, \omega_n, \chi_n(\omega, \phi_n), x_n(\omega)) \prod_{\substack{n=1, \dots, N, \\ \omega_n = M+1}} P_n^e(\omega'_n, \chi_n^e(\omega, \phi_n^e), x_n^e(\omega)),$$

where $P_n(\cdot)$ gives the probability that incumbent firm n transits from state ω_n to state ω'_n conditional on its exit decision being $\chi_n(\omega, \phi_n)$ and its investment decision being $x_n(\omega)$ and $P_n^e(\cdot)$ gives the probability that potential entrant n transits to state ω'_n . 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.

Since 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

⁷We may nest their formulation by setting $\bar{x}^e = 0$.

to obtain the probability that the industry transits from state ω to state ω' :

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

To see this, recall that scrap values and setup costs are independently distributed across firms. Since, from the point of view of other firms, the probability that incumbent firm n remains in the industry in state ω is $\xi_n(\omega) = \int \chi_n(\omega, \phi_n) dF(\phi_n)$, a particular realization $\iota = (\iota_1, \dots, \iota_N)$ of incumbent firms' exit decisions occurs with probability $\prod_{n=1, \dots, N} \xi_n(\omega)^{\iota_n} (1 - \xi_n(\omega))^{1-\iota_n}$. Similarly, a particular realization $\iota^e = (\iota_1^e, \dots, \iota_N^e)$ of potential entrants' entry decisions occurs with probability $\prod_{n=1, \dots, N} \xi_n^e(\omega)^{\iota_n^e} (1 - \xi_n^e(\omega))^{1-\iota_n^e}$. Equation (1) results from observing that if $\omega_n \neq M+1$ ($\omega_n = M+1$), then firm n is an incumbent (entrant) and conditioning on all possible realizations of incumbent firms' exit decisions ι and potential entrants' entry decisions ι^e .

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 $\xi(\omega)$ and $\xi^e(\omega)$. Thus, when forming an expectation over the industry's future state, a firm does not need to know the entire exit and entry rules $\chi_{-n}(\omega, \cdot)$ and $\chi_{-n}^e(\omega, \cdot)$ of its rivals; rather it suffices to know their exit and entry probabilities.

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 firm 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 firm that can sell off its assets for one dollar may behave very differently than an otherwise identical incumbent firm 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

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

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 while $V_n(\omega, \phi_n)$ is the value function *after* the firm has drawn its scrap value, $V_n(\omega)$ is the expected value function, i.e., the value function *before* the firm has drawn its scrap value. The RHS 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 inside 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, $\beta E \{ V_n(\omega') | \cdot \}$. Several remarks are in order. First, since 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 the other incumbents, $\xi_{-n}(\omega)$ and $x_{-n}(\omega)$, as well as on the decisions of all potential entrants, $\xi^e(\omega)$ and $x^e(\omega)$. It also conditions its expectations on its own investment choice and presumes that it remains in the industry in state ω , i.e., it conditions on $\omega'_n \neq M+1$. Note that with the recursive formulation of the incumbent's problem in equation (2) there is no need to condition on firms' entire strategies.

Since 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 $\xi_{-n}(\omega)$ and $\xi^e(\omega)$ in light of equation (1).

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

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

where

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

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, provided that the distribution of scrap values $F(\cdot)$ has a continuous and positive density, incumbent firm n has a unique optimal exit choice for all scrap values (except for a set of measure zero). 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:

1. with the cutoff scrap value $\bar{\phi}_n(\omega)$ itself; or
2. 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 $F^{-1}(\xi_n(\omega)) = \bar{\phi}_n(\omega)$.⁸ 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

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

Two essential points should be noted: First, an optimizing incumbent cares about the expectation of the scrap value conditional on collecting it, $E \left\{ \phi_n | \phi_n > F^{-1}(\tilde{\xi}_n(\omega)) \right\}$, rather than its unconditional expectation $E(\phi_n) = \phi$. The term $\int_{\phi_n > F^{-1}(\tilde{\xi}_n(\omega))} (\phi_n - \phi) dF(\phi_n) = (1 - \tilde{\xi}_n(\omega)) \left(E \left\{ \phi_n | \phi_n > F^{-1}(\tilde{\xi}_n(\omega)) \right\} - \phi \right)$ captures the difference between the conditional and the unconditional expectation. It reflects our assumption that scrap values are random and, consequently, is not present in a game of complete information such as Ericson & 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 since the

⁸If the support of $F(\cdot)$ is bounded, we define $F^{-1}(0)$ ($F^{-1}(1)$) to be the infimum (supremum) of the support.

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

$$V_n^e(\omega, \phi_n^e) = \sup_{\substack{\tilde{\chi}_n^e(\omega, \phi_n^e) \in \{0,1\}, \\ \tilde{x}_n^e(\omega, \phi_n^e) \in [0, \bar{x}^e]}} \tilde{\chi}_n^e(\omega, \phi_n^e) \left\{ -\phi_n^e - \tilde{x}_n^e(\omega, \phi_n^e) + \beta E \{ V_n(\omega') | \omega, \omega'_n \neq M + 1, \tilde{x}_n^e(\omega, \phi_n^e), \xi(\omega), \xi_{-n}^e(\omega), x(\omega), x_{-n}^e(\omega) \} \right\}. \quad (5)$$

Unlike the incumbent's value function, the entrant's value function is not defined recursively. Instead, it can be easily calculated given the incumbent's value function because the entrant is short-lived and does not solve an intertemporal maximization problem to reach its decisions.⁹ Depending on the entry decision $\chi_n^e(\omega, \phi^e)$, the RHS 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 $-\phi_n^e - \tilde{x}_n^e(\omega, \phi_n^e)$. Second, the entrant takes the net present value of its future cash flows into account. Since potential entrant n becomes incumbent firm n at the end of the setup period, this is given by $\beta E \{ V_n(\omega') | \cdot \}$. The entrant conditions its expectations on the decisions of all incumbents, $\xi(\omega)$ and $x(\omega)$ as well as on the decisions of the other entrants, $\xi_{-n}^e(\omega)$ and $x_{-n}^e(\omega)$. It also conditions its expectations on its own investment choice and presumes that it enters the industry in state ω , i.e., 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. Since investment is chosen conditional on entering the industry, the optimal investment choice $x_n^e(\omega)$ is independent of the firm's setup cost ϕ_n^e . Given its investment choice, the entrant then decides whether or not to enter the industry. The optimal entry decision is characterized by

$$\xi_n^e(\omega, \phi_n^e) = \begin{cases} 1 & \text{if } \phi_n^e \leq \bar{\phi}_n^e(\omega), \\ 0 & \text{if } \phi_n^e \geq \bar{\phi}_n^e(\omega), \end{cases}$$

where

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

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

⁹It is easy to allow for long-lived entrants by adding the recursive term $(1 - \tilde{\chi}_n^e(\omega, \phi_n^e))\beta E \{ V_n^e(\omega') | \omega, \omega'_n = M + 1, \xi(\omega), \xi_{-n}^e(\omega), x(\omega), x_{-n}^e(\omega) \}$, where $V_n^e(\omega) = \int V_n^e(\omega, \phi_n^e) dF^e(\phi_n^e)$ is the expected value function, to equation (5).

form $1[\phi_n^e < \bar{\phi}_n^e(\omega)]$. The set of all such rules can be indexed by the cutoff setup cost $\bar{\phi}_n^e(\omega)$ or by the corresponding probability $\xi_n^e(\omega)$ of potential entrant n entering the industry in state ω . Imposing the reservation property and integrating over ϕ_n^e on both sides of equation (5) yields

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

where $V_n^e(\omega) = \int V_n^e(\omega, \phi_n^e) dF^e(\phi_n^e)$ is the expected value function. The term $-\int_{\phi_n^e < F^{e-1}(\tilde{\xi}_n^e(\omega))} (\phi_n^e - \phi^e) dF^e(\phi_n^e)$ is again not present in a setting with complete information.

Notation. To save on notation, we identify the n th incumbent firm with firm n in states $\omega_n \neq M+1$ and the n th potential entrant with firm n in state $\omega_n = M+1$ in what follows. That is, we define

$$\begin{aligned} V_n^e(\omega_1, \dots, \omega_{n-1}, \omega_n, \omega_{n+1}, \dots, \omega_N) &= V_n(\omega_1, \dots, \omega_{n-1}, M+1, \omega_{n+1}, \dots, \omega_N), \\ \xi_n^e(\omega_1, \dots, \omega_{n-1}, \omega_n, \omega_{n+1}, \dots, \omega_N) &= \xi_n(\omega_1, \dots, \omega_{n-1}, M+1, \omega_{n+1}, \dots, \omega_N), \\ x_n^e(\omega_1, \dots, \omega_{n-1}, \omega_n, \omega_{n+1}, \dots, \omega_N) &= x_n(\omega_1, \dots, \omega_{n-1}, M+1, \omega_{n+1}, \dots, \omega_N). \end{aligned}$$

Let $S = \Omega^N = \{\omega^1, \dots, \omega^{|S|}\}$. Define the $|S| \times N$ matrix V by

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

and the $|S| \times (N-1)$ matrix V_{-n} by $V_{-n} = (V_1, \dots, V_{n-1}, V_{n+1}, \dots, V_N)$. 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 $V(\omega) = (V_1(\omega), \dots, V_N(\omega))$ and $V_{-n}(\omega) = (V_1(\omega), \dots, V_{n-1}(\omega), V_{n+1}(\omega), \dots, V_N(\omega))$. 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.

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

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

denotes firm n 's feasible actions in state ω . A strategy or policy for firm n 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. Define $\mathcal{U}_n = \times_{\omega \in S} \mathcal{U}_n(\omega)$ to be the strategy space of firm n and $\mathcal{U} = \times_{n=1}^N \mathcal{U}_n$ to be the strategy space of the entire industry. By construction $\mathcal{U}_n(\omega)$ and hence \mathcal{U}_n and \mathcal{U} are nonempty, convex, and compact assuming that $\bar{x} < \infty$ and $\bar{x}^e < \infty$.

Using the above notation, the Bellman equations (4) and (7) of incumbent firm n and potential entrant n , respectively, can more compactly be stated as

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

where

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

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, p. 166) or local income function (Whitt 1980, p. 35).

Equilibrium. Our solution concept is that of 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 RHS 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, since 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 Markov perfect equilibrium is subgame perfect. Note that since a best reply to Markovian strategies u_{-n} is a Markovian strategy u_n , a Markov perfect equilibrium remains a subgame perfect equilibrium even if more general strategies are considered.

3 Examples

In this section we provide a series of 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 & Pakes (1995) model, then a symmetric

equilibrium in pure entry/exit strategies may fail to exist, contrary to their assertion.¹⁰ 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. In the remainder of this section we explain how to solve the incomplete-information game using a slightly modified version of Pakes & McGuire's (1994) algorithm and we argue that this is no more demanding than computing a (possibly nonexistent) equilibrium in pure entry/exit strategies of the complete-information game whereas the alternative of computing an equilibrium in mixed entry/exit strategies is infeasible in most (if not all) applications of Ericson & Pakes's (1995) framework.

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, since 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 things further, we assume that entry is prohibitively costly and focus entirely on exit.¹¹ 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

$$\frac{\beta\pi(1, 1)}{1 - \beta} < \phi < \frac{\beta\pi(1, 2)}{1 - \beta}. \quad (10)$$

Hence, while 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:

$$\begin{aligned} V(1, 2) &= \sup_{\tilde{\chi}(1, 2) \in \{0, 1\}} \pi(1, 2) + (1 - \tilde{\chi}(1, 2))\phi + \tilde{\chi}(1, 2)\beta V(1, 2), \\ V(1, 1) &= \sup_{\tilde{\chi}(1, 1) \in \{0, 1\}} \pi(1, 1) + (1 - \tilde{\chi}(1, 1))\phi + \tilde{\chi}(1, 1)\beta \left\{ \chi(1, 1)V(1, 1) + (1 - \chi(1, 1))V(1, 2) \right\}. \end{aligned}$$

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

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

¹⁰We defer a formal definition of our symmetry notion to section 4.2.

¹¹A similar example can be constructed to demonstrate that there may not exist a symmetric equilibrium in pure entry strategies.

where

$$\bar{\phi}(1, 2) = \beta V(1, 2), \quad (11)$$

$$\bar{\phi}(1, 1) = \beta \left\{ \chi(1, 1)V(1, 1) + (1 - \chi(1, 1))V(1, 2) \right\}. \quad (12)$$

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

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

This contradicts assumption (10); 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

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

This contradicts assumption (10); 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

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

which again contradicts assumption (10). Hence, there cannot be a symmetric equilibrium in pure exit strategies.¹²

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

$$\begin{aligned} V(1, 2) &= \frac{\pi(1, 2)}{1 - \beta}, \quad V(1, 1) = \pi(1, 1) + \phi, \\ \xi(1, 2) &= 1, \quad \xi(1, 1) = \frac{(1 - \beta)\phi - \beta\pi(1, 2)}{\beta((1 - \beta)(\pi(1, 1) + \phi) - \pi(1, 2))}. \end{aligned}$$

Example: Random scrap values/setup costs. Pakes & 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 duopolis-

¹²In this particular example there exist two asymmetric equilibria in pure exit strategies. In each of them, one firm exits in state (1, 1), the other remains in the industry.

tic industry and prescribing that both firms should exit.

To restore existence we introduced in section 2 random scrap values in addition to the random setup costs suggested by Pakes & McGuire (1994). We now modify the above example to illustrate this use of incomplete information. Specifically, we assume that scrap values are independently and identically distributed across firms and periods, and that its scrap value is private to a firm. 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 $E(\theta) = 0$. The Bellman equation of firm 1 is

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

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 $\tilde{\xi}(\omega) = F\left(\frac{\bar{\phi}(\omega) - \phi}{\epsilon}\right)$,¹³ where

$$\begin{aligned} \bar{\phi}(1, 2) &= \beta V(1, 2), \\ \bar{\phi}(1, 1) &= \beta \left\{ \xi(1, 1)V(1, 1) + (1 - \xi(1, 1))V(1, 2) \right\}. \end{aligned}$$

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 $\xi(1, 1)$.

Obtaining analytic solutions is complicated by the fact that the equations that define the value function are no longer linear in $V(\omega)$ because $V(\omega)$ enters $\bar{\phi}(\omega)$. For analytic convenience, let θ be uniformly distributed on the interval $[-1, 1]$. This implies

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

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 parameters 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 $\epsilon \leq 5$, it involves $\xi(1, 2) = 1$ and $0 < \xi(1, 1) < 1$. Table 2 describes the equilibrium for various values of ϵ . Given the parameter values in Table

¹³To see this, note that the first and second derivatives of the RHS of the Bellman equation are given by $\frac{d(\cdot)}{d\xi(\omega)} = -\phi - \epsilon F^{-1}(\tilde{\xi}(\omega)) + \bar{\phi}(\omega)$ and $\frac{d^2(\cdot)}{d\xi(\omega)^2} = -\epsilon \frac{1}{F'(F^{-1}(\xi(\omega)))}$, respectively.

| parameter | $\pi(1, 1)$ | $\pi(1, 2)$ | ϕ | β |
|-----------|-------------|-------------|--------|-----------------|
| value | 0 | 1 | 15 | $\frac{20}{21}$ |

Table 1: Parameter 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 |

Table 2: Equilibrium with random scrap values.

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 the next section, we show that existence and convergence are general properties of the game of incomplete information.

Computational issues. The advantage of studying a game of incomplete information is that it eliminates the need for mixed entry/exit strategies without jeopardizing existence. While the state space is effectively the same as in the Ericson & Pakes (1995) model, the Bellman equation in the game of incomplete information has one additional term reflecting how random scrap values/setup costs affect a firm's per-period payoff. But, as demonstrated in the above example, an appropriate distribution of the scrap values/setup costs yields a closed-form expression for this term.¹⁴ Thus introducing incomplete information into the Ericson & Pakes (1995) model adds essentially nothing to the computational burden and it ensures that the search for an equilibrium is never hopeless because an equilibrium in cutoff entry/exit strategies always exists. Below we first show that a slightly modified version of the algorithm that Pakes & McGuire (1994) developed can be used to compute equilibria of the game of incomplete information. Then we review the alternative of computing mixed-strategy equilibria of the game of complete information and argue that the existing algorithms are incapable of solving problems as large as the ones that arise in applications of Ericson & Pakes's (1995) framework.

Like Pakes & McGuire's (1994) algorithm, the modified algorithm works iteratively.

¹⁴Besides the uniform distribution many others yield a closed-form expression, including triangular and Beta distributions.

| ϵ | $\lambda = 1$ | $\lambda = 0.1$ | $\lambda = 0.01$ | $\lambda = 0.001$ | $\lambda = 0.0001$ |
|------------|---------------|-----------------|------------------|-------------------|--------------------|
| 10 | 87 | 457 | 3620 | 28416 | 207741 |
| 5 | 251 | 1256 | 9294 | 67105 | 421664 |
| 1 | | 325 | 1610 | 13641 | 113229 |
| 0.1 | | | 1555 | 13092 | 107742 |
| 0.01 | | | | 13092 | 107742 |
| 0.001 | | | | | 107742 |

Table 3: Number of iterations until convergence.

In the context of the above example the l th iteration takes a value function V^l and a policy function ξ^l as its input and outputs updated value and policy functions V^{l+1} and ξ^{l+1} . Each iteration proceeds as follows: First, update the policy function by assigning $\xi^{l+1}(\omega) \leftarrow F\left(\frac{\bar{\phi}(\omega) - \phi}{\epsilon}\right)$, where

$$\bar{\phi}(1, 2) = \beta V^l(1, 2), \quad (13)$$

$$\bar{\phi}(1, 1) = \beta \left\{ \xi^l(1, 1) V^l(1, 1) + (1 - \xi^l(1, 1)) V^l(1, 2) \right\}. \quad (14)$$

Second, update the value function by assigning

$$V^{l+1}(\omega) \leftarrow \pi(\omega) + (1 - \xi^{l+1}(\omega))\phi + \epsilon \int_{\theta > F^{-1}(\xi^{l+1}(\omega))} \theta dF(\theta) + \xi^{l+1}(\omega) \bar{\phi}(\omega),$$

where $\bar{\phi}(\omega)$ is as in equations (13) and (14) except that $\xi^l(\omega)$ is replaced by $\xi^{l+1}(\omega)$. The algorithm terminates once the relative change in the value and the policy functions from one iteration to the next is below a pre-specified tolerance. We take this tolerance to be 10^{-8} and use $V^0 = 0$ and $\xi^0 = 0$ as starting values.

In the column labelled $\lambda = 1$ Table 3 lists the number of iterations until convergence. The algorithm converges quickly if ϵ is large but fails to converge otherwise (indicated by a blank). It turns out that adding a dampening scheme (see e.g. Chapter 3 of Judd 1998) aids convergence. The dampening scheme combines the updated and the current policy function with the assignment $\xi^{l+1}(\omega) \leftarrow \lambda F\left(\frac{\bar{\phi}(\omega) - \phi}{\epsilon}\right) + (1 - \lambda)\xi^l(\omega)$ where $\lambda \in (0, 1)$. The remaining columns of Table 3 list the number of iterations until convergence for different values of $\lambda \in (0, 1)$. Roughly speaking, we are able to decrease ϵ by an order of magnitude if we are willing to do the same with λ . This results in a tenfold increase in the number of iterations.

Three conclusions emerge. First, the modified algorithm succeeds in computing an equilibrium in cutoff entry/exit strategies. In sharp contrast, the original algorithm fails in cases in which only an equilibrium in (nondegenerate) mixed strategies exists, for it can neither exactly compute nor closely approximate such an equilibrium. Second, if random scrap values/setup costs are deemed to be an accurate description of industry fundamentals, then the modified algorithm converges quickly to an equilibrium in cutoff entry/exit strategies

even if there is limited variation in the scrap values/setup costs. In the context of the above example $\epsilon = 5$ ($\epsilon = 1$) implies that the scrap value lies within $\pm 33\%$ ($\pm 7\%$) of its mean.¹⁵ Third, if, on the other hand, scrap values/setup costs are thought to be deterministic, then the game of incomplete information may be useful to obtain an approximation to an equilibrium in mixed entry/exit strategies of the game of complete information, although it appears that the computational burden increases with the required accuracy.

The alternative to introducing incomplete information into the Ericson & Pakes (1995) model is computing mixed-strategy equilibria of the game of complete information. In his survey of the literature on dynamic stochastic games with discrete actions Breton (1991) laments: “In the zero-sum case, there exist reasonably efficient algorithms, but such is not the case in the general sum N -player case” (p. 56). Using a mathematical programming approach he reports being able to solve, with considerable difficulty, dynamic stochastic games with up to 3 players, 5 states, and 5 actions per player and state. Most recently, Herings & Peeters (2004) have solved games with up to 5 players, 5 states, and 5 actions per player and state. In its simplest form the Ericson & Pakes (1995) model reduces to a dynamic entry/exit game. If there are N homogeneous firms, then the state of the industry is defined by the number of active firms and $|\Omega| = N + 1$. The state of the art algorithms for dynamic stochastic games are thus able to compute an equilibrium in mixed entry/exit strategies in industries with a handful of firms. If, on the other hand, firms are heterogeneous so that profits from product market competition depend on the identity of the active firms, then the state of the industry indicates for each firm whether it is active or inactive. Hence, $|\Omega| = 2^N$ and computing an equilibrium in mixed entry/exit strategies in industries with three firms pushes the limits of the best algorithms.

Of course, the real strength of Ericson & Pakes’s (1995) framework is that it allows us to model dynamic competition between heterogeneous firms that decide on investment in addition to entry and exit, thereby capturing the empirical finding that firm heterogeneity evolves endogenously in response to random occurrences in the investment process. This greatly increases the computational burden and puts most (if not all) applications out of reach of the above algorithms. First, these algorithms are designed for discrete actions whereas investment is continuous. The obvious solution is to discretize investment. This, however, adds to the number of actions per firm and state, and even a fairly coarse discretization easily renders the problem intractable. Second, any sensible specification of endogenously evolving firm heterogeneity requires more than just distinguishing between active and inactive firms. In fact, all payoff-relevant characteristics of a firm have to be encoded in its state. Consequently, the smallest applications of Ericson & Pakes’s (1995) framework have hundreds and the largest ones millions of states (despite symmetry). The sheer size of the state space alone makes these applications orders of magnitude too large to be solvable with the best algorithms for computing mixed-strategy equilibria in dynamic

¹⁵In his study of the airframe industry Benkard (2004) assumes that the scrap value (setup cost) lies within $\pm 40\%$ ($\pm 17\%$) of its mean.

stochastic games.

In contrast, the algorithms suggested by Pakes & McGuire (1994, 2001) are capable of solving such large games. The reason is that they are built around iterated best reply ideas. That is, in a given state, the algorithm updates the decisions of one firm at a time (holding the decisions of its rivals' fixed), then it moves on to the next state. Best replies are easy to calculate, often in closed-form. While this conveys a massive advantage to these algorithms, the drawback is that they are confined to pure-strategy equilibria (which may not exist).¹⁶ Our use of incomplete information combines the best of both worlds, i.e., the computational tractability of pure strategies with the theoretical soundness of mixed strategies. Indeed, as we have argued above, incomplete information does not add to the computational burden and yet it guarantees that an equilibrium in cutoff entry/exit strategies always exists. The next section formally establishes this claim.

4 Existence and Convergence

In this section, we show how incorporating firm heterogeneity in the form of random scrap values/setup costs into the Ericson & Pakes (1995) model guarantees the existence of an equilibrium. We first 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 papers use models that are sufficiently abstract to enable us to construct the bulk of the existence proof by citing their intermediate results. We then build on our basic existence result in three ways. We first show that a symmetric equilibrium exists. Requiring the 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 by-product, this last result implies that there exists an equilibrium in the Ericson & Pakes (1995) model provided that mixed entry/exit strategies are admissible.

¹⁶A fruitful venue for future research may be to abandon these iterated best reply ideas in favor of computing the Nash equilibrium in one state at a time. The idea is that, in each state, firms play a static game taking their value functions as given. This suggests an algorithm consisting of an outer loop that searches for a fixed point in value functions and an inner loop that computes the Nash equilibrium in each state. Govindan & Wilson (2003, 2004) compute mixed-strategy Nash equilibria in normal form games with up to 6 players and 16 actions per player and up to 12 players and 14 actions per player, respectively, thus leaving room to discretize investment, and report running times ranging from a few seconds to a few minutes. Once these running times are multiplied by the number of states (inner loop) and the number of iterations (outer loop), however, it is unclear whether the suggested algorithm is practical.

4.1 Existence

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, i.e., $N < \infty$ and $M < \infty$. (ii) Profits are bounded, i.e., there exists $\bar{\pi} < \infty$ such that $-\bar{\pi} < \pi_n(\omega) < \bar{\pi}$ for all ω and all n . (iii) Investments are bounded, i.e., $\bar{x} < \infty$ and $\bar{x}^e < \infty$. (iv) The distributions of scrap values $F(\cdot)$ and setup costs $F^e(\cdot)$ have continuous and positive densities and bounded supports, i.e., there exist $\bar{\phi} < \infty$ and $\bar{\phi}^e < \infty$ such that the supports of $F(\cdot)$ and $F^e(\cdot)$ are contained in the interval $[-\bar{\phi}, \bar{\phi}]$ and $[-\bar{\phi}^e, \bar{\phi}^e]$, respectively. (v) Firms discount future payoffs, i.e., $\beta \in [0, 1)$.

Next we assume continuity of firm n 's local income function $h_n(\cdot)$. Similar continuity assumptions are commonplace in the literature on dynamic stochastic games (see Mertens 2002).

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 .

Note that $h_n(\cdot)$ is always continuous in V_n as long as V_n enters $h_n(\cdot)$ via the expected value of firm n 's future cash flows, $E\{V_n(\omega')|\cdot\}$. Moreover, given that in our model formulation current profit is additively separable from investment, continuity of $h_n(\cdot)$ merely requires continuity of the transition function $P(\cdot)$. We make the continuity assumption on $h_n(\cdot)$ rather than on $P(\cdot)$ to facilitate the adaptation of our existence proof to other models in which current profit is not additively separable from investment.¹⁷

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, computing mixed strategies over continuous actions is not practical. To guarantee the existence of an equilibrium in cutoff entry/exit and 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 .¹⁸

¹⁷In models of learning-by-doing (Cabral & Riordan 1994, Benkard 2004, Besanko, Doraszelski, Kryukov & Satterthwaite 2007), for example, firms' price or quantity decisions today determine their current profit as well as their marginal cost of production tomorrow. Hence, the current profit of incumbent firm n is $\pi_n(\omega, x(\omega))$, where $x(\omega) = (x_1(\omega), \dots, x_N(\omega))$ denotes the prices charged or the quantities marketed.

¹⁸Assumption 3 can be weakened to hold for all possible maximal return functions $V_{n,u_{-n}} \in [\underline{V}^*, \bar{V}^*]^{|S|}$ rather than for all possible value functions V_n , where the loose upper and lower bounds are given by $\bar{V}^* = \bar{\phi}^e + \frac{\bar{\pi}}{1-\beta} + \bar{\phi}$ and $\underline{V}^* = -\bar{\phi}^e - \frac{\bar{\pi}}{1-\beta} - \bar{\phi}$.

In section 5 we define UIC admissibility of the transition function $P(\cdot)$ and prove that this condition on the model's primitives ensures uniqueness of investment choice and, thus, existence of an equilibrium that is amenable to computation.

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 $h_n(\omega, \xi_n(\omega), x_n(\omega), u_{-n}(\omega), V_n)$ 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 although 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 1 *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.¹⁹ Fix strategies u_{-n} and consider firm n 's problem. Since 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. From a computational point of view, the proof mimics an algorithm that nests a dynamic programming problem within a fixed point problem.²⁰

Before stating the proof of proposition 1, we introduce and discuss a number of constructs that will also be useful in later parts of the paper. We start with the decision problem. Let \mathcal{V}_n denote the space of bounded $|S| \times 1$ vectors with the sup norm and let ρ denote the corresponding metric. Fix $u_{-n} \in \mathcal{U}_{-n}$ and define the *maximal return operator* $H_{n,u_{-n}}^* : \mathcal{V}_n \rightarrow \mathcal{V}_n$ 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

¹⁹Given that standard arguments (e.g., Federgruen 1978, Whitt 1980) 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 1 for the proofs of propositions 3, 4, and 5.

²⁰Such an algorithm has indeed been suggested by Rust (1994).

described by V_n . Note that the RHS of the above equation coincides with the RHS of the Bellman equation (9).

Since profits, investments, scrap values, and setup costs are bounded by assumption 1, $H_{n,u-n}^*$ takes bounded vectors into bounded vectors. Application of Blackwell's sufficient conditions (monotonicity and discounting, see e.g. p. 54 of Stokey & Lucas (1989)) shows that $H_{n,u-n}^*$ is a contraction with modulus β . The contraction mapping theorem (Stokey & Lucas 1989, p. 50) therefore 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,

$$V_{n,u-n}^*(\omega) = \sup_{\tilde{u}_n(\omega) \in \mathcal{U}_n(\omega)} h_n(\omega, \tilde{u}_n(\omega), u_{-n}(\omega), V_{n,u-n}^*) \quad (15)$$

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

$$V_{n,u}(\omega) = h_n(\omega, u(\omega), V_{n,u}) \quad (16)$$

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

Note that there is a tight connection between the return function $V_{n,u}$ and the maximal return function $V_{n,u-n}^*$. In fact, because the return operator $H_{n,u}$ is monotonic, theorem 3 of Denardo (1967) establishes that

$$V_{n,u-n}^*(\omega) = \sup_{\tilde{u}_n \in \mathcal{U}_n} V_{n,\tilde{u}_n,u-n}(\omega) \quad (17)$$

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}) . Put loosely, choosing one's optimal response state by state yields the same return as choosing one's optimal response jointly for all states. Somewhat more formally, the solution to the Bellman equation (9) coincides with the solution to the (considerably more cumbersome) sequence form of the decision problem. To bring out the implications of equa-

tion (17), fix strategies u_{-n} and consider a family of games where each member of the family is indexed by the initial state ω . Firm n 's best reply to u_{-n} for the game beginning in state ω yields a payoff of $\sup_{\tilde{u}_n \in \mathcal{U}_n} V_{n,\tilde{u}_n,u_{-n}}(\omega)$. But the maximal return function $V_{n,u_{-n}}^*$ is independent of the initial state and so is the strategy defined pointwise by $\arg \sup_{\tilde{u}_n(\omega) \in \mathcal{U}_n(\omega)} h_n(\omega, \tilde{u}_n(\omega), u_{-n}(\omega), V_{n,u_{-n}}^*)$. Hence, the best reply can always be taken to be independent of the initial state, a fact which we shall use presently.

With this machinery in place, we turn to the game problem. Consider the mapping $\Upsilon : \mathcal{U} \rightarrow \mathcal{U}$ defined pointwise by

$$\Upsilon_n(u) = \left\{ \tilde{u}_n \in \mathcal{U}_n : \tilde{u}_n(\omega) \in \arg \sup_{\tilde{u}_n(\omega) \in \mathcal{U}_n(\omega)} h_n(\omega, \tilde{u}_n(\omega), u_{-n}(\omega), V_{n,u_{-n}}^*) \text{ for all } \omega \right\}. \quad (18)$$

Note that $\Upsilon_n(\cdot)$ is the best-reply correspondence of firm n . 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 1. We begin by establishing that $\Upsilon(\cdot)$ is non-empty and upper hemi-continuous. 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 . Assumption 2 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 $\Upsilon_n(\cdot)$ 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).

His lemma 3.2 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 . His lemma 3.1 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 (17) 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} .

Since $h_n(\omega, u_n(\omega), u_{-n}(\omega), V_{n,u_{-n}}^*)$ 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 e.g. p. 62 of Stokey & Lucas 1989) 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 non-empty and upper hemi-continuous in u_{-n} . Since ω was arbitrary, this establishes that $\Upsilon_n(\cdot)$ is a non-empty and upper hemi-continuous correspondence that maps \mathcal{U}_{-n} into \mathcal{U}_n . Hence, $\Upsilon(\cdot) = (\Upsilon_1(\cdot), \dots, \Upsilon_N(\cdot))$ is non-empty and upper hemi-continuous.

We next show that $\Upsilon(\cdot)$ is single-valued. Fix ω . 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. 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

²¹We set $\mathcal{W}_n = \mathcal{V}_n$ to obtain a special case of Whitt's (1980) lemma.

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

Since $\Upsilon(\cdot)$ is non-empty, single-valued, and upper hemi-continuous, it is, in fact, a continuous function that maps the non-empty, convex, and compact set \mathcal{U} into itself. Brouwer's fixed point theorem therefore applies: a $u \in \mathcal{U}$ exists such that $u \in \Upsilon(u)$. ■

4.2 Symmetry

In section 4.1 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.

Informally the notion of symmetry in Ericson & 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, then 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, then 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 $\kappa = (\kappa_1, \dots, \kappa_N)$ be a permutation of $(1, \dots, N)$. The policy functions $u = (u_1, \dots, u_N)$ are symmetric if

$$u_n(\omega_{\kappa_1}, \dots, \omega_{\kappa_{n-1}}, \omega_{\kappa_n}, \omega_{\kappa_{n+1}}, \dots, \omega_{\kappa_N}) = u_{\kappa_n}(\omega_1, \dots, \omega_{n-1}, \omega_n, \omega_{n+1}, \dots, \omega_N) \quad (19)$$

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 $V = (V_1, \dots, V_N)$ are symmetric and satisfy the analog of equation (19).

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

1. 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 (19) to obtain $u_2(3, 3, 2) = u_2(\omega_2, \omega_3, \omega_1) = u_3(\omega_1, \omega_2, \omega_3) = u_3(2, 3, 3)$.
2. 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 (19) 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 & Pakes (1995): A firm's policy is a 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 $\kappa = (n, 2, \dots, n-1, 1, n+1, \dots, N)$ in equation (19) to obtain

$$u_n(\omega_n, \omega_2, \dots, \omega_{n-1}, \omega_1, \omega_{n+1}, \dots, \omega_N) = u_1(\omega_1, \omega_2, \dots, \omega_{n-1}, \omega_n, \omega_{n+1}, \dots, \omega_N), \quad (20)$$

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 $\kappa = (1, 2, \dots, k-1, l, k+1, \dots, l-1, k, l+1, \dots, N)$ with $k \geq 2$ and $l \geq 2$ in equation (19) to obtain

$$u_1(\omega_1, \omega_2, \dots, \omega_l, \dots, \omega_k, \dots, \omega_N) = u_1(\omega_1, \omega_2, \dots, \omega_k, \dots, \omega_l, \dots, \omega_N), \quad (21)$$

and similarly for the value function. That is, only the firm's own state and the distribution of rivals' states matters. The latter property is commonly referred to as anonymity or as exchangeability.²²

We are now ready to state our symmetry assumption.

Assumption 4 *The local income functions are symmetric, i.e.,*

$$h_n(\omega_{\kappa_1}, \dots, \omega_{\kappa_N}, u_{\kappa_1}(\omega), \dots, u_{\kappa_N}(\omega), V_n) = h_{\kappa_n}(\omega_1, \dots, \omega_N, u_1(\omega), \dots, u_N(\omega), V_{\kappa_n}) \quad (22)$$

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 to the principle that identical actions in identical situations yield identical payoffs. Let $n = 2$ and $\kappa = (2, 3, 1)$ in equation (22) 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 states ω_1 and one in state ω_2 . Since 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.

²²Equations (20) and (21) are often together taken as the definition of symmetry (e.g., Doraszelski & Pakes 2007). It is easy to see that they are equivalent to our notion of symmetry in equation (19). Working with equation (19) instead of equations (20) and (21) simplifies the notation in the remainder of this section.

While 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.

Proposition 2 *Assumption 4 is satisfied provided that (i) the profit functions are symmetric, i.e.,*

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

for all ω , n and all κ and (ii) the transition function is symmetric, i.e.,

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

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

The proof of proposition 2 is straightforward but tedious and therefore omitted. Note that in the special case of independent transitions, condition (ii) of proposition 2 is satisfied whenever the factors $P_n(\cdot)$ of the transition function $P(\cdot)$ are the same across firms, i.e., $P_n(\omega'_n, \omega_n, \chi_n(\omega, \phi_n), x_n(\omega)) = P_1(\omega'_1, \omega_1, \chi_1(\omega, \phi_1), x_1(\omega))$ for all n .

Together with assumptions 1, 2, and 3 in section 4.1, assumption 4 ensures existence of a symmetric equilibrium.

Proposition 3 *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 the Ericson-Pakes 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 1 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 3 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 $\sigma_n = (\sigma_{n,1}, \dots, \sigma_{n,M}, \sigma_{n,M+1})$, where $\sigma_{n,m}$ denotes the number of competitors of firm n that are in state m (excluding firm n), and $\Sigma = \left\{ \sigma_n \in \{0, 1, \dots, N-1\}^{M+1} \mid \sum_{m=1}^{M+1} \sigma_{n,m} = N-1 \right\}$ 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.²³

Define the inverse function $\tau_n^{-1} : S^\circ \rightarrow S$ 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 $\omega_1 \leq \omega_2 \leq \dots \leq \omega_{n-1} \leq \omega_{n+1} \leq \dots \leq \omega_N$. 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, $(3, 2, 2, 5) = \tau_1^{-1}(\tau_1(3, 2, 5, 2))$, $(2, 2, 3, 5) = \tau_2^{-1}(\tau_2(3, 2, 5, 2))$, etc. A state $\tilde{\omega}$ is called *canonical* if and only if $\tilde{\omega} = \tau_1^{-1}(\tilde{\omega}_1, \sigma_1)$ for some $(\tilde{\omega}_1, \sigma_1)$. We use the symbol \sim 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, i.e., $u_n^\circ(\omega_n, \sigma_n) = u_1^\circ(\omega_n, \sigma_n)$ 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

$$\begin{aligned} & h_1^\circ((\omega_1, \sigma_1), u_1^\circ(\omega_1, \sigma_1), u_2^\circ(\tau_2(\tau_1^{-1}(\omega_1, \sigma_1))), \dots, u_N^\circ(\tau_N(\tau_1^{-1}(\omega_1, \sigma_1))), V_1^\circ) \\ & = h_1(\tau_1^{-1}(\omega_1, \sigma_1), u_1^\circ(\omega_1, \sigma_1), u_2^\circ(\tau_2(\tau_1^{-1}(\omega_1, \sigma_1))), \dots, u_N^\circ(\tau_N(\tau_1^{-1}(\omega_1, \sigma_1))), \Lambda_1(V_1^\circ)), \end{aligned} \quad (23)$$

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 $H_{1, u_1^\circ}^* : \mathcal{V}_1^\circ \rightarrow \mathcal{V}_1^\circ$

²³Gowrisankaran (1999b) develops an algorithm for the efficient representation of the reduced state space.

pointwise by

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

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 1. 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 $\Upsilon_1^\circ : \mathcal{U}_1^\circ \rightarrow \mathcal{U}_1^\circ$ by

$$\Upsilon_1^\circ(u_1^\circ) = \left\{ \tilde{u}_1^\circ \in \mathcal{U}_1^\circ : \tilde{u}_1^\circ(\omega_1, \sigma_1) \in \arg \sup_{\tilde{u}_1^\circ(\omega_1, \sigma_1) \in \mathcal{U}_1^\circ(\omega_1, \sigma_1)} h_1^\circ((\omega_1, \sigma_1), \tilde{u}_1^\circ(\omega_1, \sigma_1), u_1^\circ(\tau_2(\tau_1^{-1}(\omega_1, \sigma_1))), \dots, u_1^\circ(\tau_N(\tau_1^{-1}(\omega_1, \sigma_1))), V_{1,u_1^\circ}^{\circ*}) \text{ for all } (\omega_1, \sigma_1) \right\}. \quad (24)$$

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 1, $\Upsilon_1^\circ(\cdot)$ is non-empty, single-valued, and upper hemi-continuous and thus a function to which Brouwer's fixed point theorem applies.

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

$$u_n = \Lambda_n(u_1^\circ). \quad (25)$$

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

$$V_{n,u_n}^* = \Lambda_n(V_{1,u_1^\circ}^{\circ*}). \quad (26)$$

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. This justifies the common practice of computing equilibria directly on the reduced state space.

Proof of proposition 3. 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 n th 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

$$\begin{aligned} & h_n(\omega, u_1(\omega), u_2(\omega), \dots, u_{n-1}(\omega), \tilde{u}_n(\omega), u_{n+1}(\omega), \dots, u_N(\omega), V_{n,u_n}^*) \\ &= h_1(\hat{\omega}, \tilde{u}_n(\omega), u_2(\omega), \dots, u_{n-1}(\omega), u_1(\omega), u_{n+1}(\omega), \dots, u_N(\omega), V_{1,u_{-1}}^*) \\ &= h_1(\hat{\omega}, \tilde{u}_n(\omega), u_2(\hat{\omega}), \dots, u_{n-1}(\hat{\omega}), u_n(\hat{\omega}), u_{n+1}(\hat{\omega}), \dots, u_N(\hat{\omega}), V_{1,u_{-1}}^*), \end{aligned}$$

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) non-canonical 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 $\kappa_{-1} = (\kappa_2, \dots, \kappa_N)$ of $(2, \dots, N)$. We have $\mathcal{U}_1(\hat{\omega}) = \mathcal{U}_1(\check{\omega})$ for the set of feasible actions and, for an arbitrary action $\tilde{u}_1(\hat{\omega}) \in \mathcal{U}_1(\hat{\omega})$,

$$\begin{aligned} & h_1(\hat{\omega}, \tilde{u}_1(\hat{\omega}), u_2(\hat{\omega}), \dots, u_N(\hat{\omega}), V_{1,u_{-1}}^*) \\ &= h_1(\check{\omega}, \tilde{u}_1(\hat{\omega}), u_{\kappa_2}(\hat{\omega}), \dots, u_{\kappa_N}(\hat{\omega}), V_{1,u_{-1}}^*) \\ &= h_1(\check{\omega}, \tilde{u}_1(\hat{\omega}), u_2(\check{\omega}), \dots, u_N(\check{\omega}), V_{1,u_{-1}}^*), \end{aligned}$$

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

$$\begin{aligned} & h_1(\check{\omega}, \tilde{u}_1(\check{\omega}), u_2(\check{\omega}), \dots, u_N(\check{\omega}), V_{1,u_{-1}}^*) \\ &= h_1(\check{\omega}, \tilde{u}_1(\check{\omega}), u_1^\circ(\tau_2(\check{\omega})), \dots, u_1^\circ(\tau_N(\check{\omega})), \Lambda_1(V_{1,u_1^\circ}^*)) \\ &= h_1(\tau_1^{-1}(\check{\omega}_1, \sigma_1), \tilde{u}_1(\check{\omega}), u_1^\circ(\tau_2(\tau_1^{-1}(\check{\omega}_1, \sigma_1))), \dots, u_1^\circ(\tau_N(\tau_1^{-1}(\check{\omega}_1, \sigma_1))), \Lambda_1(V_{1,u_1^\circ}^*)) \\ &= h_1^\circ((\check{\omega}_1, \sigma_1), \tilde{u}_1(\check{\omega}), u_1^\circ(\tau_2(\tau_1^{-1}(\check{\omega}_1, \sigma_1))), \dots, u_1^\circ(\tau_N(\tau_1^{-1}(\check{\omega}_1, \sigma_1))), V_{1,u_1^\circ}^*), \end{aligned} \quad (27)$$

where the first equality follows from equations (25) and (26), the second from the fact that $\check{\omega} = \tau^{-1}(\check{\omega}_1, \sigma_1)$ for some $(\check{\omega}_1, \sigma_1)$ because $\check{\omega}$ is canonical, and the last from equation (23). Moreover, we have $\mathcal{U}_1(\check{\omega}) = \mathcal{U}_1(\tau_1^{-1}(\check{\omega}_1, \sigma_1)) = \mathcal{U}_1^\circ(\check{\omega}_1, \sigma_1)$ for the set of feasible actions. Since the last line of equation (27) is the maximand of firm 1 in the best-reply correspondence in equation (24), firm 1 has no incentive to deviate. Moreover, since the problem of firm n in state ω is identical to the problem of firm 1 in state $\check{\omega}$ by the first two steps of the proof, no firm has an incentive to deviate from the candidate equilibrium. ■

4.3 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 $E(\theta_n) = E(\theta_n^e) = 0$.

Firm n 's return or local income function $h_n^\epsilon(\cdot)$ becomes

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

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

Proposition 1 in section 4.1 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 we have already pointed out, there is a need to allow for mixed entry/exit strategies in a game with deterministic scrap values/setup costs such as Ericson & Pakes (1995). We thus ask if 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 4 *Suppose assumptions 1, 2, and 3 hold and consider a sequence $\{\epsilon^l\}$ such that $\lim_{l \rightarrow \infty} \epsilon^l = 0$. Let $\{u^l\}$ be a corresponding sequence of equilibria in cutoff entry/exit strategies such that $\lim_{l \rightarrow \infty} u^l = u$. Then u is an equilibrium in mixed entry/exit strategies.*

Proof. Let $\{V_{n,u^l}^{\epsilon^l}\}$ 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 (16) in section 4.1 shows that each element of $\{V_{n,u^l}^{\epsilon^l}\}$ is well-defined due to assumption 1. Moreover, since $H_{n,u}^\epsilon V_n$ is continuous in ϵ and u for all V_n , lemma 3.2 of Whitt (1980) implies that the return function $V_{n,u}^\epsilon$ is continuous in ϵ and u . Let $V_{n,u} = \lim_{l \rightarrow \infty} V_{n,u^l}^{\epsilon^l}$ for all n .

The proof proceeds in two steps. In the first step, we verify that the limiting strategies u_n are 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

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

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

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

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 . Since $V_{n,u}$ is a fixed point of the maximal return operator of the game of complete information, it is the maximal return function. ■

Note that proposition 4 does not imply that $\lim_{l \rightarrow \infty} u^l$ exists. On the other hand, since \mathcal{U} is compact every sequence $\{u^l\}$ has a convergent subsequence, and proposition 4 applies to the subsequential limit. This establishes

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

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 . This is restrictive because, in general, the value function V_n of firm n and the actions $u_{-n}(\omega)$ of its rivals may take on values such that $h_n(\dots, 1, x_n(\omega), \dots)$ attains its maximum at more than one investment level. To see the role assumption 3 plays, suppose for the moment that it is violated. Then $\Upsilon_n(\cdot)$, the best-reply correspondence of firm n , is no longer guaranteed to be a function, thus necessitating the use of Kakutani's instead of Brouwer's fixed point theorem. Kakutani's fixed point theorem, in turn, requires convex-valuedness of $\Upsilon_n(\cdot)$. Using standard arguments, convex-valuedness can be ensured by allowing for mixed investment strategies. This, however, is not practical because computing mixed strategies over continuous actions is well beyond present computational capabilities.

Fortunately, a judicious choice of transition probabilities guarantees that the investment choice is unique. In this section we define unique investment choice (UIC) admissibility of the transition function $P(\cdot)$ and show in proposition 5 that if this condition on the model's primitives is satisfied, then an equilibrium in cutoff entry/exit and pure investment

strategies exists. We then give 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*

$$K_n(\omega', \omega, \chi_{-n}(\omega, \phi_{-n}), x_{-n}(\omega)) Q_n(\omega, x_n(\omega)) + L_n(\omega', \omega, \chi_{-n}(\omega, \phi_{-n}), x_{-n}(\omega)), \quad (28)$$

where $Q_n(\omega, x_n(\omega))$ is twice differentiable, strictly increasing, and strictly concave in $x_n(\omega)$, i.e.,

$$\frac{d}{dx_n(\omega)} Q_n(\omega, x_n(\omega)) > 0, \quad \frac{d^2}{dx_n(\omega)^2} Q_n(\omega, x_n(\omega)) < 0 \quad (29)$$

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).²⁴

UIC admissibility ensures that firm n 's local income function $h_n(\dots, 1, x_n(\omega), \dots)$ 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.²⁵ We are now ready to state our main result establishing that a computationally tractable equilibrium exists in our model.

Proposition 5 *Suppose assumptions 1 and 2 hold. If the transition function $P(\cdot)$ is UIC admissible, then an equilibrium exists in cutoff entry/exit and pure investment strategies. If in addition assumption 4 holds, then a symmetric equilibrium exists in cutoff entry/exit and pure investment strategies.*

Proof. In light of propositions 1 and 3, it suffices to show that assumption 3 holds. Since 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, $E\{V_n(\omega') | \omega, \omega'_n \neq M+1, \xi_{-n}(\omega), x(\omega)\}$, in its local income function $h_n(\dots, 1, x_n(\omega), \dots)$ can be written in a separable form as

$$A_n(\omega, u_{-n}(\omega), V_n) Q_n(\omega, x_n(\omega)) + B_n(\omega, u_{-n}(\omega), V_n). \quad (30)$$

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

²⁴Condition 1 can be generalized to allow for $Q(\cdot)$ to depend on $x_{-n}(\omega)$.

²⁵Of course, uniqueness of investment choice can also be achieved by other means. In particular, if $\bar{x}^e = 0$, then a potential entrant has no choice but to invest zero, thereby stripping the potential entrant of any control over its initial position within the industry (as in Ericson & Pakes 1995).

moves from state ω to state ω' . Hence,

$$\begin{aligned}
& \sum_{\omega' \in S} V_n(\omega') \sum_{\iota_{-n} \in \{0,1\}^{N-1}} P(\omega', \omega, 1, \iota_{-n}, x(\omega)) \prod_{k \neq n} \xi_k(\omega)^{\iota_k} (1 - \xi_k(\omega))^{1-\iota_k} \\
&= \sum_{\omega' \in S} V_n(\omega') \sum_{\iota_{-n} \in \{0,1\}^{N-1}} \left[K_n(\omega', \omega, \iota_{-n}, x_{-n}(\omega)) Q_n(\omega, x_n(\omega)) + L_n(\omega', \omega, \iota_{-n}, x_{-n}(\omega)) \right] \\
&\quad \times \prod_{k \neq n} \xi_k(\omega)^{\iota_k} (1 - \xi_k(\omega))^{1-\iota_k} \\
&= \underbrace{\left[\sum_{\omega' \in S} V_n(\omega') \sum_{\iota_{-n} \in \{0,1\}^{N-1}} K_n(\omega', \omega, \iota_{-n}, x_{-n}(\omega)) \prod_{k \neq n} \xi_k(\omega)^{\iota_k} (1 - \xi_k(\omega))^{1-\iota_k} \right]}_{A_n(\omega, u_{-n}(\omega), V_n)} Q_n(\omega, x_n(\omega)) \\
&\quad + \underbrace{\left[\sum_{\omega' \in S} V_n(\omega') \sum_{\iota_{-n} \in \{0,1\}^{N-1}} L_n(\omega', \omega, \iota_{-n}, x_{-n}(\omega)) \prod_{k \neq n} \xi_k(\omega)^{\iota_k} (1 - \xi_k(\omega))^{1-\iota_k} \right]}_{B_n(\omega, u_{-n}(\omega), V_n)}
\end{aligned}$$

where the first equality uses the separability condition (28).

Next we differentiate $h_n(\dots, 1, x_n(\omega), \dots)$ with respect to $x_n(\omega)$. By virtue of equation (30), the FOC 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 FOC 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 (29). Hence, $x_n(\omega) = \hat{x}_n(\omega)$ is the unique maximizer. If there does not exist a solution to the FOC 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 $x_n(\omega) = 0$. ■

UIC admissibility allows for much more flexibility in the transition probabilities than the simple schemes seen in the extant literature where each firm is restricted to each period 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 (see Besanko & Doraszelski 2004) where a firm's state describes its capacity. In each period, the firm decides how much to spend on an investment

| | $\omega'_2 = \omega_2 + 1$ | $\omega'_2 = \omega_2$ | $\omega'_2 = \omega_2 - 1$ |
|----------------------------|--|---|---|
| $\omega'_1 = \omega_1 + 1$ | $(1 - \delta)p_1(1 - \delta)p_2$ | $(1 - \delta)p_1[\delta p_2 + (1 - \delta)(1 - p_2)]$ | $(1 - \delta)p_1\delta p_2$ |
| $\omega'_1 = \omega_1$ | $[\delta p_1 + (1 - \delta)(1 - p_1)]$ $\times (1 - \delta)p_2$ | $[\delta p_1 + (1 - \delta)(1 - p_1)]$ $\times [\delta p_2 + (1 - \delta)(1 - p_2)]$ | $[\delta p_1 + (1 - \delta)(1 - p_1)]$ $\times \delta p_2$ |
| $\omega'_1 = \omega_1 - 1$ | $\delta(1 - p_1)(1 - \delta)p_2$ | $\delta(1 - p_1)[\delta p_2 + (1 - \delta)(1 - p_2)]$ | $\delta(1 - p_1)\delta p_2$ |

Table 4: Transition probabilities. Independent transitions to immediately adjacent states.

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 $\omega \in \{2, \dots, M - 1\}^2$ are given in Table 4.²⁶

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

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

This expression satisfies the separability condition (28), as do the corresponding expressions for the probabilities of moving to some other state $\omega' \neq \omega$. In addition, the derivative condition (29) 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 & 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 $\omega \in \{2, \dots, M - 1\}^2$ are given in Table 5.

²⁶Edge states must be treated specially. If $\omega_n = 1$, then the probability of moving up to state $\omega'_n = 2$ (remaining in state $\omega'_n = 1$) is $(1 - \delta)p_n$ (δp_n); if $\omega_n = M$, then the probability of dropping down to state $\omega'_n = M - 1$ (remaining in state $\omega'_n = M$) is $\delta(1 - p_n)$ ($1 - \delta(1 - p_n)$).

| | $\omega'_2 = \omega_2 + 1$ | $\omega'_2 = \omega_2$ | $\omega'_2 = \omega_2 - 1$ |
|----------------------------|------------------------------|---|------------------------------|
| $\omega'_1 = \omega_1 + 1$ | $(1 - \delta) p_1 p_2$ | $(1 - \delta) p_1 (1 - p_2)$ | 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 | $\delta (1 - p_1) p_2$ | $\delta (1 - p_1) (1 - p_2)$ |

Table 5: Transition probabilities. Dependent transitions to immediately adjacent states.

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 (28) and (29) 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 firm states. For example, the probability that firm n moves from its initial state ω_n to the intermediate state $\hat{\omega}_n \in \{1, \dots, M\}$ may be

$$\left\{ \begin{array}{lll} \zeta_{n, \omega_n, 1} + \eta_{n, \omega_n, 1} p_n & \text{if} & \hat{\omega}_n = 1, \\ \vdots & \vdots & \vdots \\ \zeta_{n, \omega_n, \omega_n - 1} + \eta_{n, \omega_n, \omega_n - 1} p_n & \text{if} & \hat{\omega}_n = \omega_n - 1, \\ \zeta_{n, \omega_n, \omega_n} + \eta_{n, \omega_n, \omega_n} p_n & \text{if} & \hat{\omega}_n = \omega_n, \\ \zeta_{n, \omega_n, \omega_n + 1} + \eta_{n, \omega_n, \omega_n + 1} p_n & \text{if} & \hat{\omega}_n = \omega_n + 1, \\ \vdots & \vdots & \vdots \\ \zeta_{n, \omega_n, M} + \eta_{n, \omega_n, M} p_n & \text{if} & \hat{\omega}_n = M, \end{array} \right.$$

where $x_n(\omega)$ affects the probability of a transition from state ω_n to state $\hat{\omega}_n$ either positively or negatively depending on the sign of $\eta_{n, \omega_n, \hat{\omega}_n}$.²⁷ 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 (29). For example, let

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

²⁷The 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 $\sum_{\hat{\omega}_n=1}^M \eta_{n, \omega_n, \hat{\omega}_n} = 0$.

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, while 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, e.g., by the inter-quartile range.

6 Multiplicity

It is widely believed that multiplicity is not an issue in Ericson & Pakes's (1995) framework. Pakes (2000) summarizes his experience as follows:

... we have experimented quite a bit with the core version of the algorithm, and we never found two sets of equilibrium policies for a given set of primitives (we frequently run the algorithm several times using different initial conditions or different orderings of points looking for other equilibria that might exist). We should emphasize here that the core version, and indeed most other versions that have been used, all use quite simple functional forms for the primitives of the problem, and multiplicity of equilibrium may well be more likely when more complicated functional forms are used. Of course, most applied work suffices with quite simple functional forms. (pp. 18–19)

In this section we discuss three examples that show that there need not be a unique equilibrium that is symmetric. Our goal here is solely to establish the possibility of multiple equilibria. Throughout we use the “quite simple functional forms” alluded to by Pakes (2000). Using more complicated functional forms in a model of learning-by-doing, Besanko et al. (2006) find as many as nine equilibria for some parameterizations and discuss the economics underlying this multiplicity.

Example: Investment decisions. We build on the game of capacity accumulation from section 5. There are $N = 2$ firms with $M \geq 3$ “active” states. In state ω_n firm n 's capacity is \bar{q}_{ω_n} . Transitions are limited to immediately adjacent states and are independent across firms.²⁸ Products are undifferentiated and firms compete in prices subject to capacity

²⁸Because the transition function $P(\cdot)$ is UIC admissible, it is guaranteed that multiplicity is not due to a violation of assumption 3.

| parameter | M | \bar{q}_1 | \bar{q}_2 | \dots | \bar{q}_{10} | v | m | α | δ | β |
|-----------|-----|-------------|-------------|---------|----------------|-----|-----|----------|----------|-----------------|
| value | 10 | 0 | 5 | \dots | 45 | 1 | 40 | 2.375 | 0.03 | $\frac{20}{21}$ |

Table 6: Parameter values.




Figure 1: Two equilibria.

constraints. There are m identical consumers with unit demand and reservation price v . The equilibrium of this Bertrand-Edgeworth product market game is characterized in Chapter 2 of Ghemawat (1997). Let $\pi(\omega_1, \omega_2)$ denote firm 1's current profit in state $\omega = (\omega_1, \omega_2)$. Symmetry implies that firm 2's current profit in state ω is $\pi(\omega_2, \omega_1)$. Table 6 gives the parameter values.

Figure 1 illustrates the value and policy functions of two symmetric equilibria. In both investment activity is greatest in states on or near the diagonal of the state space. That is, firms with equal or similar capacities are engaged in a preemption race to become the industry leader. The difference in investment activity is greatest in state $(5, 5)$ where both firms invest 1.90 in the first equilibrium compared to 1.03 in the second one. Investment activity also differs considerably in states $(1, 6)$ and $(6, 1)$: in the first (second) equilibrium the smaller firm invests 2.24 (3.92) and the larger firm invests 1.57 (1.46). That is, in the second equilibrium, the laggard invests heavily in a bid to catch up with the leader, and the leader to some extent accommodates the laggard. Note that multiplicity in this example rests on the dynamic nature of the game. Because product market competition takes place before investment decisions are carried out, a firm has no incentive to invest if $\beta = 0$. Hence,

multiple equilibria cannot possibly arise in the static version of the game.

The computations were performed using a Matlab 5.3 implementation of the Pakes & McGuire (1994) algorithm. The first equilibrium was computed using a Gauss-Jacobi scheme to update the value and policy functions, the second using a Gauss-Seidel scheme (see e.g. Judd 1998). This is worth pointing out because many applications of Ericson & Pakes's (1995) framework have searched for multiple equilibria by selecting a single algorithm and varying the starting values. This approach, however, failed to identify the different equilibria in our example, and its use may thus lead one to falsely conclude that multiplicity is not an issue.

Example: Entry/exit decisions. In the above example nonuniqueness results solely from firms' investment decisions in a model without entry and exit. We next show that nonuniqueness may also result from firms' exit decisions. We slightly extend our example with random scrap values/setup costs from section 3. In particular, suppose that each firm can now be in one of two "active" states (i.e., $M = 2$) and that the current profit in states $(1, 1)$, $(1, 2)$, $(2, 1)$, and $(2, 2)$ is the same. Suppose finally that a firm cannot transit between its active states (in the above example this corresponds to a situation with ineffective investment ($\alpha = 0$) and zero depreciation ($\delta = 0$)).

Again there are two symmetric equilibria. One equilibrium has both firms play the cutoff exit strategies from section 3 in states $(1, 1)$, $(1, 2)$, $(2, 1)$, and $(2, 2)$ and remain in the industry with high probability ($\xi(1, 1) = \xi(1, 2) = \xi(2, 1) = \xi(2, 2) = 0.8549$). The top panels of Table 7 summarize this equilibrium for $\epsilon = 1$. Another equilibrium is illustrated in the bottom panels of Table 7. In states $(1, 1)$ and $(2, 2)$ both firms continue to play the cutoff exit strategies from section 3 and remain in the industry with high probability ($\xi(1, 1) = \xi(2, 2) = 0.8549$). In state $(1, 2)$ firm 1 stays with low probability ($\xi(1, 2) = 0.1542$) and firm 2 stays for sure ($\xi(2, 1) = 1$) whereas in state $(2, 1)$ firm 1 stays for sure ($\xi(2, 1) = 1$) and firm 2 stays with low probability ($\xi(1, 2) = 0.1542$). Note that this equilibrium is symmetric by construction.

The two equilibria differ starkly from each other: In the first equilibrium, a duopolistic industry may over time turn into either a monopolistic or an empty industry. In the second equilibrium, if the industry starts in states $(1, 2)$ or $(2, 1)$, then it always ends up as a monopoly.

Pakes & McGuire (1994) have previously conjectured that nonuniqueness may result from firms' exit decisions, especially in "situations in which one firm out of two (or several firms out of many) must exit, but either exiting would generate an equilibrium" (p. 570). To resolve this issue they suggest to assume that firms make exit decisions sequentially in a given order. The drawback is that ordering firms renders symmetry inapplicable (since we must now distinguish between firms with the same state according to their order), thereby increasing the size of the state space and the computational burden. Further, Pakes & McGuire's (1994) suggestion may fail to rule out all but one equilibrium. To see this,

| $V(\omega_1, \omega_2)$ | $\omega_2 = 1$ | $\omega_2 = 2$ | $\omega_2 = 3$ | $\xi(\omega_1, \omega_2)$ | $\omega_2 = 1$ | $\omega_2 = 2$ | $\omega_2 = 3$ |
|-------------------------|----------------|----------------|----------------|---------------------------|----------------|----------------|----------------|
| $\omega_1 = 1$ | 15.7309 | 15.7309 | 21 | $\omega_1 = 1$ | 0.8549 | 0.8549 | 1 |
| $\omega_1 = 2$ | 15.7309 | 15.7309 | 21 | $\omega_1 = 2$ | 0.8549 | 0.8549 | 1 |
| $V(\omega_1, \omega_2)$ | $\omega_2 = 1$ | $\omega_2 = 2$ | $\omega_2 = 3$ | $\xi(\omega_1, \omega_2)$ | $\omega_2 = 1$ | $\omega_2 = 2$ | $\omega_2 = 3$ |
| $\omega_1 = 1$ | 15.7309 | 15.0238 | 21 | $\omega_1 = 1$ | 0.8549 | 0.1542 | 1 |
| $\omega_1 = 2$ | 19.8279 | 15.7309 | 21 | $\omega_1 = 2$ | 1 | 0.8549 | 1 |

Table 7: Two more equilibria.

consider our example with deterministic scrap values/setup costs from section 3. Suppose that there is just one “active” state (i.e., $M = 1$) and that, in each period, firm 1 makes its exit decision before firm 2. Hence, firm 1 takes into account how its exit decision affects that of firm 2. It is easy to see that there are two equilibria. One equilibrium has firm 1 stay in the industry in state $(1, 1)$ and firm 2 exit whereas another equilibrium has firm 1 exit and firm 2 stay. Intuitively, there is little reason to believe that sequential exit decisions alleviate the multiplicity problem unless this assumption enables one to apply a backwards induction argument. But this is not the case in a model with an infinite horizon such as the Ericson & Pakes (1995) model.

Example: Product market competition. We close this section by noting that we treat the profit function $\pi_n(\cdot)$ as a primitive. Instead we could have gone back to demand and cost fundamentals and explicitly modelled competition in the product market. To the extent that this game admits more than one equilibrium $\pi_n(\cdot)$ fails to be determined uniquely, thereby making product market competition yet another source of multiplicity.

7 Conclusions

This paper 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 & 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 & 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

Since computing mixed strategies over continuous actions is well beyond present computational capabilities, it is vital to ensure existence of an equilibrium in pure investment strategies in addition to cutoff entry/exit strategies. We achieve this in our proofs 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 several ways. We first show that a symmetric equilibrium exists provided the model's primitives are symmetric. This is a significant result for two reasons. First, from a computational viewpoint, symmetry is needed to control the size of the state space. Second, from a substantive viewpoint, in models of dynamic competition with entry and exit, there is often no compelling reason why a particular entrant should be different from any other entrant. This makes a symmetric equilibrium an especially compelling solution concept because, in such an equilibrium, firm heterogeneity arises endogenously from the idiosyncratic outcomes that the *ex ante* identical firms realize from their investments. To our knowledge, this is the first attempt to guarantee existence of a symmetric equilibrium in a broad class of dynamic stochastic games.^{29,30} Our argument is readily extended to arbitrary dynamic stochastic games.

Next 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. Our argument again extends to arbitrary dynamic stochastic games with discrete actions. While one suspects that there are fundamental similarities between the set of Markov perfect equilibria in these games and the set of Nash equilibria in normal form games, proof is still required and by no means obvious.³¹ Indeed, we have been unable to determine whether or not the approachability part of Harsanyi's (1973a) purification theorem carries over from static games to dynamic stochastic games. That is, are all equilibria of the original game approached by some equilibrium of the perturbed game as the perturbation vanishes? We leave this as an open question for future research.

²⁹Dutta & Sundaram (1992) show that there exists a symmetric Markov perfect equilibrium in two-player resource extraction games with a one-dimensional state space.

³⁰While Pesendorfer & Schmidt-Dengler (2003) purport to show the existence of a symmetric Markov perfect equilibrium, their (unproven) corollary 1 is flawed. First, their notion of symmetry is ill-conceived. In the case of two players, for example, it implies that the per-period payoffs to both players are equal, $\pi_1(a_1, a_2, s_1, s_2) = \pi_2(a_1, a_2, s_1, s_2)$, irrespective of their actions, a_1 and a_2 , and their states, s_1 and s_2 , whereas symmetry is properly defined as $\pi_1(a_1, a_2, s_1, s_2) = \pi_2(a_2, a_1, s_2, s_1)$ (see section 4.2). Second, they place no restrictions on the transition function $g(a, s, s')$. However, if the actions of different players have different effects on the future state of the game, then a symmetric equilibrium can hardly be expected to exist.

³¹For example, the counterpart to the generic finiteness and oddness of the set of Nash equilibria in normal form games (Wilson 1971, Harsanyi 1973b) has only recently been established by Haller & Lagunoff (2000) and Herings & Peeters (2004).

Finally, we provide the first example of multiple symmetric equilibria in the literature initiated by Ericson & Pakes (1995). While this formally settles the uniqueness issue, it is just an initial step. In fact, little is known to date about uniqueness of equilibrium in dynamic stochastic games. Haller & Lagunoff (2000) show that the number of equilibria is generically finite and Amir (2002) shows that there exists a unique equilibrium in pure strategies in *finite* horizon games that satisfy monotonicity, supermodularity, and dominant-diagonal conditions. More research along these lines is clearly needed.

References

- Aguirregabiria, V. & Mira, P. (2007), ‘Sequential estimation of dynamic discrete games’, *Econometrica* **75**(1), 1–54.
- Amir, R. (2002), ‘Complementarities and diagonal dominance in discounted stochastic games’, *Annals of Operations Research* **114**(1), 39–56.
- Bajari, P., Benkard, L. & Levin, J. (2007), ‘Estimating dynamic models of imperfect competition’, *Econometrica* **75**(5), 1331–1370.
- Benkard, L. (2004), ‘A dynamic analysis of the market for wide-bodied commercial aircraft’, *Review of Economic Studies* **71**(3), 581–611.
- Beresteanu, A. & Ellickson, P. (2005), The dynamics of retail oligopoly, Working paper, Duke University, Durham.
- Berry, S. & Pakes, A. (1993), ‘Some applications and limitations of recent advances in empirical industrial organization: Merger analysis’, *American Economic Review* **83**(2), 247–252.
- Besanko, D. & Doraszelski, U. (2004), ‘Capacity dynamics and endogenous asymmetries in firm size’, *Rand Journal of Economics* **35**(1), 23–49.
- Besanko, D., Doraszelski, U., Kryukov, Y. & Satterthwaite, M. (2007), Learning-by-doing, organizational forgetting, and industry dynamics, Working paper, Harvard University, Cambridge.
- Besanko, D., Doraszelski, U., Lu, L. & Satterthwaite, M. (2006), Capacity investment dynamics in oligopolistic industries, Working paper, Northwestern University, Evanston.
- Breton, M. (1991), Algorithms for stochastic games, *in* T. Raghavan, T. Ferguson, T. Parthasarathy & O. Vrieze, eds, ‘Stochastic games and related topics: In honor of Professor L.S. Shapley’, Kluwer Academic Publishers, Dordrecht, pp. 45–57.
- Cabral, L. & Riordan, M. (1994), ‘The learning curve, market dominance, and predatory pricing’, *Econometrica* **62**(5), 1115–1140.
- Chakrabarti, S. (2003), ‘Pure strategy Markov equilibrium in stochastic games with a continuum of players’, *Journal of Mathematical Economics* **39**, 693–724.
- Collard-Wexler, A. (2005), Plant turnover and demand fluctuations in the ready-mix concrete industry, Working paper, New York University, New York.

- Curtat, L. (1996), ‘Markov equilibria of stochastic games with complementarities’, *Games and Economic Behavior* **17**(2), 177–199.
- Denardo, E. (1967), ‘Contraction mappings in the theory underlying dynamic programming’, *SIAM Review* **9**(2), 165–177.
- Doraszelski, U. & Pakes, A. (2007), A framework for applied dynamic analysis in IO, in M. Armstrong & R. Porter, eds, ‘Handbook of Industrial Organization’, Vol. 3, North-Holland, Amsterdam, pp. 1887–1966.
- Dutta, P. & Sundaram, R. (1992), ‘Markovian equilibrium in a class of stochastic games: Existence theorems for discounted and undiscounted models’, *Economic Theory* **2**, 197–214.
- Erdem, E. & Tybout, J. (2003), Trade policy and industrial sector responses: Using evolutionary models to interpret the evidence, Working paper no. 9947, NBER, Cambridge.
- Ericson, R. & Pakes, A. (1995), ‘Markov-perfect industry dynamics: A framework for empirical work’, *Review of Economic Studies* **62**, 53–82.
- Escobar, J. (2007), Existence of pure and behavior strategy stationary Markov equilibrium in dynamic stochastic games, Working paper, Stanford University, Stanford.
- Federgruen, A. (1978), ‘On N -person stochastic games with denumerable state space’, *Advances in Applied Probability* **10**, 452–471.
- Fershtman, C. & Pakes, A. (2000), ‘A dynamic oligopoly with collusion and price wars’, *Rand Journal of Economics* **31**, 294–326.
- Fink, A. (1964), ‘Equilibrium points of stochastic noncooperative games’, *Journal of Science Hiroshima University Series A* **28**, 89–93.
- Fudenberg, D., Gilbert, R., Stiglitz, J. & Tirole, J. (1983), ‘Preemption, leapfrogging and competition in patent races’, *European Economic Review* **22**, 3–31.
- Ghemawat, P. (1997), *Games businesses play*, MIT Press, Cambridge.
- Goettler, R., Parlour, C. & Rajan, U. (2005), ‘Equilibrium in a dynamic limit order market’, *Journal of Finance* **60**(5), 2149–2192.
- Govindan, S. & Wilson, R. (2003), ‘A global Newton method to compute Nash equilibria’, *Journal of Economic Theory* **110**(1), 65–86.
- Govindan, S. & Wilson, R. (2004), ‘Computing Nash equilibria by iterated polymatrix approximation’, *Journal of Economic Dynamics and Control* **28**, 1229–1241.
- Gowrisankaran, G. (1995), A dynamic analysis of mergers, Ph.D. thesis, Yale University, New Haven.
- Gowrisankaran, G. (1999a), ‘A dynamic model of endogenous horizontal mergers’, *Rand Journal of Economics* **30**(1), 56–83.
- Gowrisankaran, G. (1999b), ‘Efficient representation of state spaces for some dynamic models’, *Journal of Economic Dynamics and Control* **23**, 1077–1098.

- Gowrisankaran, G. & Town, R. (1997), ‘Dynamic equilibrium in the hospital industry’, *Journal of Economics and Management Strategy* **6**(1), 45–74.
- Haller, H. & Lagunoff, R. (2000), ‘Genericity and Markovian behavior in stochastic games’, *Econometrica* **68**(5), 1231–1248.
- Harris, C. & Vickers, J. (1987), ‘Racing with uncertainty’, *Review of Economic Studies* **54**(1), 1–21.
- Harsanyi, J. (1973*a*), ‘Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points’, *International Journal of Game Theory* **2**(1), 1–23.
- Harsanyi, J. (1973*b*), ‘Oddness of the number of equilibrium points: A new proof’, *International Journal of Game Theory* **2**(4), 235–250.
- Herings, J. & Peeters, R. (2004), ‘Stationary equilibria in stochastic games: Structure, selection, and computation’, *Journal of Economic Theory* **118**, 32–60.
- Jenkins, M., Liu, P., Matzkin, R. & McFadden, D. (2004), The browser war: Econometric analysis of Markov perfect equilibrium in markets with network effects, Working paper, Stanford University, Stanford.
- Judd, K. (1998), *Numerical methods in economics*, MIT Press, Cambridge.
- Kadyrzhanova, D. (2006), The leader-bias hypothesis: Corporate control dynamics in industry equilibrium, Working paper, Columbia University, New York.
- Maskin, E. & Tirole, J. (2001), ‘Markov perfect equilibrium, I: Observable actions’, *Journal of Economic Theory* **100**, 191–219.
- McKelvey, R. & McLennan, A. (1996), Computation of equilibria in finite games, in H. Amman, D. Kendrick & J. Rust, eds, ‘Handbook of Computational Economics’, North-Holland, Amsterdam, pp. 87–142.
- Mertens, J. (2002), Stochastic games, in R. Aumann & S. Hart, eds, ‘Handbook of Game Theory’, Vol. 3, Elsevier, Amsterdam, pp. 1810–1832.
- Pakes, A. (2000), A framework for applied dynamic analysis in I.O., Working paper no. 8024, NBER, Cambridge.
- Pakes, A. & McGuire, P. (1994), ‘Computing Markov-perfect Nash equilibria: Numerical implications of a dynamic differentiated product model’, *Rand Journal of Economics* **25**(4), 555–589.
- Pakes, A. & McGuire, P. (2001), ‘Stochastic algorithms, symmetric Markov perfect equilibrium, and the “curse” of dimensionality’, *Econometrica* **69**(5), 1261–1281.
- Pakes, A., Ostrovsky, M. & Berry, S. (2006), ‘Simple estimators for the parameters of discrete dynamic games (with entry/exit examples)’, *Rand Journal of Economics* **forthcoming**.
- Pesendorfer, M. & Schmidt-Dengler, P. (2003), Identification and estimation of dynamic games, Working paper no. 9726, NBER, Cambridge.

- Rust, J. (1994), Estimation of dynamic structural models, problems and prospects: discrete decision processes, *in* C. Sims, ed., ‘Advances in econometrics: Sixth World Congress’, Vol. 2, Cambridge University Press, Cambridge.
- Ryan, S. (2005), The costs of environmental regulation in a concentrated industry, Working paper, Duke University, Durham.
- Seim, K. (2001), Spatial differentiation and market structure: The video rental industry, Ph.D. thesis, Yale University, New Haven.
- Sobel, M. (1971), ‘Noncooperative stochastic games’, *The Annals of Mathematical Statistics* **42**(6), 1930–1935.
- Stokey, N. & Lucas, R. with Prescott, E. (1989), *Recursive methods in economic dynamics*, Harvard University Press, Cambridge.
- Whitt, W. (1980), ‘Representation and approximation of noncooperative sequential games’, *SIAM Journal of Control and Optimization* **18**(1), 33–48.
- Wilson, R. (1971), ‘Computing equilibria of n-person games’, *SIAM Journal of Applied Mathematics* **21**, 80–87.