

Doraszelski, Ulrich; Satterthwaite, Mark

Working Paper

Foundations of Markov-perfect industry dynamics: existence, purification, and multiplicity

Discussion Paper, No. 1383

Provided in Cooperation with:

Kellogg School of Management - Center for Mathematical Studies in Economics and Management Science, Northwestern University

Suggested Citation: Doraszelski, Ulrich; Satterthwaite, Mark (2003) : Foundations of Markov-perfect industry dynamics: existence, purification, and multiplicity, Discussion Paper, No. 1383, Northwestern University, Kellogg School of Management, Center for Mathematical Studies in Economics and Management Science, Evanston, IL

This Version is available at:

<https://hdl.handle.net/10419/31190>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

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

Ulrich Doraszelski
Hoover Institution, Stanford University[†]

Mark Satterthwaite
Kellogg School of Management, Northwestern University[‡]

November 5, 2003

Abstract

In this paper we show that existence of a Markov perfect equilibrium (MPE) in the Ericson & Pakes (1995) model of dynamic competition in an oligopolistic industry with investment, entry, and exit requires admissibility of mixed entry/exit strategies, contrary to Ericson & Pakes's (1995) assertion. This is problematic because the existing algorithms cannot cope with mixed strategies. To establish a firm basis for computing dynamic industry equilibria, we 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 a MPE in cutoff entry/exit strategies and is computationally no more demanding than the original game of complete information. Building on our basic existence result, we first show that a symmetric and anonymous MPE 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, MPEs in cutoff entry/exit strategies converge to MPEs in mixed entry/exit strategies of the game of complete information. Next, we provide a condition on the model's primitives that ensures the existence of a MPE in pure investment strategies. Finally, we provide the first example of multiple symmetric and anonymous MPEs in this literature.

1 Introduction

The empirical literature on industry dynamics has established two key findings (Dunne, Roberts & Samuelson 1988). First, entry and exit occur simultaneously. Second, there is heterogeneity among firms, and this heterogeneity evolves endogenously in response to random occurrences, for example, in the investment process. To capture these findings,

*Thanks to Michaela Draganska, Dino Gerardi, Ken Judd, Patricia Langohr, and Eilon Solan. Satterthwaite acknowledges gratefully that this material is based upon work supported by the National Science Foundation under Grant No. 0121541.

[†]Stanford, CA 94305-6010, U.S.A., doraszelski@hoover.stanford.edu.

[‡]Evanston, IL 60208-2009, U.S.A., m-satterthwaite@kellogg.northwestern.edu.

Ericson & Pakes (1995) develop a general model that tracks an oligopolistic industry over time. In each period, incumbent firms decide 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. Firm heterogeneity is accounted for by encoding all payoff-relevant characteristics of a firm in its “state.” For example, a firm’s state may describe its production capacity, cost structure, or the quality of its product. A firm is able to change its state over time through investment, and the variability in the fortunes of seemingly similar firms is captured by an idiosyncratic shock that affects a firm’s transition from one state to another. Since the Ericson & Pakes (1995) model is by far too complex to be solved analytically, Pakes & McGuire (1994, 2001) provide algorithms to compute a pure-strategy Markov perfect equilibrium (MPE) of this dynamic stochastic game.

Although the framework laid out by Ericson & Pakes (1995) is attractive and has been widely applied (see Pakes (2000) for a survey), its theoretical foundations are wanting. Due to a shortcoming in Ericson & Pakes’s (1995) argument, it is unknown whether or not a MPE exists and, if it exists, whether or not it is unique. Since any attempt to compute a nonexistent equilibrium is doomed, resolving the existence issue is clearly critical. However, given that the purpose of Ericson & Pakes’s (1995) framework is to provide a computable model of dynamic industry equilibrium, two additional and equally important issues arise. First, computing mixed strategies over discrete actions such as entry and exit in dynamic stochastic games poses a formidable challenge despite the considerable progress that has been made in the context of finite games (see McKelvey & McLennan (1996)). Moreover, computing mixed strategies over continuous actions such as investment is infeasible at present. It is thus vital to guarantee existence of a MPE in pure strategies. 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 focus attention on symmetric and anonymous MPEs. If no such MPE exists, then this is a fruitless approach.

This paper resolves all the above issues: We establish that a symmetric and anonymous MPE in pure strategies always exists under reasonable conditions. We furthermore demonstrate that the MPE is not necessarily unique. These results are tailored to the specifics of the Ericson & Pakes (1995) model and fulfill the goal of providing a firm basis for computing dynamic industry equilibria using existing algorithms—notably Pakes & McGuire (1994, 2001).

There are three difficulties in devising a computationally tractable model. First, the existence of an equilibrium cannot be ensured without allowing firms to randomize, in some way or another, over discrete actions. To eliminate the need for mixed entry/exit strategies without jeopardizing existence, we generalize an idea that Pakes & McGuire (1994) suggested to overcome convergence problems in their algorithm: treating a potential entrant’s

setup cost as a random variable. In accordance with Harsanyi's (1973) technique for purifying mixed-strategy equilibria, we 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 a MPE in cutoff entry/exit strategies which can be handled by existing algorithms. 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 were following a mixed strategy. Harsanyi's (1973) insight that a perturbed game of incomplete information can purify the mixed-strategy equilibria of an underlying game of complete information enables us to settle the first difficulty in devising a computationally tractable model.

The second difficulty is to ensure pure investment strategies. The extant literature (see Mertens (2002) for a survey) routinely allows for randomization over continuous actions. Computing mixed strategies over continuous actions, however, is not practical. We therefore have to make sure that a firm's optimal investment level is always unique, for that guarantees that the MPE 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, as we discuss below, is not unduly restrictive on the transition functions that it admits.

Our paper contributes to a growing literature that establishes existence of a MPE 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. Bergin & Bernhardt (1995) and Chakrabarti (2003) analyze dynamic stochastic games with a continuum of players. Chakrabarti (2003) shows that there exists a MPE in pure strategies in such a game provided that the per-period payoffs and the transition density function depend only on the *average* response of the players. Our approach complements this literature by providing a different, and very intuitive, sufficient condition for existence of a MPE in pure strategies in a broad class of dynamic stochastic games.

The third and final difficulty in devising a computationally tractable model is to en-

sure that the MPE is not only in pure strategies, but also symmetric and anonymous. We show that this is the case under appropriate assumptions on the model's primitives. Symmetry and anonymity are important because they ease the computational burden considerably. Instead of having to compute value functions (i.e., payoffs) and policy functions (i.e., strategies) for *all* firms, under symmetry and anonymity it suffices to compute value and policy functions for *one* firm. In addition, symmetry and anonymity reduce the size of the state space on which these functions are defined. Besides its computational advantages, a symmetric and anonymous MPE 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.

With these difficulties tamed, our existence proof is straightforward. For the most part we adapt the argument of Whitt (1980) to a setting with incomplete information. While the literature provides several existence theorems for dynamic stochastic games with a finite state space and a continuum of actions (e.g., Federgruen 1978), they all allow for mixed strategies in establishing existence of a MPE that is neither guaranteed to be symmetric nor anonymous. We are unable to simply cite these existence theorems because a MPE must not involve mixed strategies and it must be symmetric and anonymous in order to be suitable for computation.

In addition to introducing random scrap values/setup costs, our model relaxes two restrictive features of the entry process specified by Ericson & Pakes (1995). First, while more than one incumbent firm can exit the industry per period, at most one potential entrant can enter it. Second, an entrant is randomly assigned to an arbitrary position and thus has no control over its initial position within in the industry. These features are especially troublesome because industry evolution frequently takes the form of a preemption race (Besanko & Doraszelski 2002, Doraszelski & Markovich 2003, Langohr 2003). 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, and an imposed order of entry may prove to be decisive for the structure of the industry. Our specification of the entry process does not suffer from this drawback. By assuming that entry decisions, like exit decisions, are made simultaneously, we allow more than one firm to enter the industry per period. Moreover, we allow an entrant to make an initial investment in order to improve the odds that it enters the industry in a more favorable state. Taken together, these changes make 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.

The plan of the paper is as follows. In Section 2, we develop the model. In Section 3 we provide a series of simple examples to illustrate the key themes of the subsequent analysis. Our first example shows that if scrap values/setup costs are constant across firms and periods as in Ericson & Pakes (1995), then a symmetric MPE in pure entry/exit strategies may fail to exist, contrary to the assertion of Ericson & Pakes (1995). Our second example shows how the incorporation of random scrap values/setup costs ensures that a symmetric MPE in cutoff entry/exit strategies exists and that this MPE converges to the MPE in mixed entry/exit strategies of the original game as the randomness in the scrap values/setup costs vanishes. In the remainder of Section 3 we argue that the incomplete-information game is computationally no more demanding than the complete-information game and demonstrate how it can be solved using a slightly modified version of Pakes & McGuire’s (1994) algorithm.

In Section 4, we turn to full-fledged incomplete-information game with investment in addition to entry and exit. In Section 4.1 we establish from first principles that there is a MPE in cutoff entry/exit and pure investment strategies. While such a MPE is computable, the burden may be substantial. In Section 4.2 we therefore adapt our basic existence proof to show existence of a symmetric and anonymous MPE in cutoff entry/exit and pure investment strategies. In Section 4.3 we show that, as the distribution of the random scrap values/setup costs becomes degenerate, a MPE in cutoff entry/exit and pure investment strategies of the incomplete-information game converges to a MPE in mixed entry/exit and pure investment strategies of the complete-information game. This immediately implies that there exists a MPE in the Ericson & Pakes (1995) model provided that mixed entry/exit strategies are admissible. Moreover, to the extent that incomplete information is viewed as a “computational trick” rather than an accurate description of industry fundamentals, the addition of random scrap values/setup costs does not change the nature of strategic interactions among firms.

In Section 5 we define UIC admissibility of the transition function and prove that this condition on the model’s primitives is sufficient for the uniqueness of a firm’s investment choice. It thus guarantees existence of a MPE in pure investment strategies. 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 exactly how to specify more general UIC admissible transition functions.

To our knowledge, all applications of Ericson & Pakes’s (1995) framework have found a single MPE. We settle the uniqueness issue in Section 6 by providing examples of multiple symmetric and anonymous MPEs. Whereas Pakes & McGuire (1994) conjecture that nonuniqueness might result from firms’ entry/exit decisions, our main example does not rely on this. Rather, nonuniqueness results solely from firms’ investment decisions. 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.¹ At any point in time the industry is completely characterized by the list of firms’ states $\omega = (\omega_1, \dots, \omega_N) \in S$ where $S = \Omega^N$ is the state space.² 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.³ 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.

¹This 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.

²Time-varying characteristics of the competitive environment are easily added to the description of the industry. Besanko & Doraszelski (2002), 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.

³Limiting 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.

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 $E(\phi_n) = \phi$.⁴ Scrap values are independently and identically distributed across firms and periods. Incumbent firm n learns its scrap value ϕ_n prior to making its exit and investment decisions, but the scrap values of its rivals remain unknown to it. Let $\chi_n(\omega, \phi_n) = 1$ 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.

⁴It is straightforward to allow firm n 's scrap value ϕ_n to vary systematically with its state ω_n by replacing $F(\cdot)$ by $F_{\omega_n}(\cdot)$.

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 $\mathbb{E}(\phi_n^e) = \phi^e$. 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. Ericson & Pakes (1995) assume that, unlike exit decisions, entry decisions are made sequentially. We propose a simultaneous formulation of entry that allows 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 being randomly assigned to an arbitrary position and thus has no control over its initial position within the industry.⁵

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 $x(\omega) =$

⁵We nest their formulation by setting $\bar{x}^e = 0$.

$(x_1(\omega), \dots, x_N(\omega))$, etc. Necessarily $\sum_{\omega' \in S} P(\omega', \omega, \chi(\omega, \phi), \chi^e(\omega, \phi^e), x(\omega), x^e(\omega)) = 1$.

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

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

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 . It 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. 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 charac-

terized 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

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

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 $F^{-1}(\xi_n(\omega)) = \bar{\phi}_n(\omega)$.⁶ 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 (2) yields

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

The term $\int_{\phi_n > F^{-1}(\tilde{\xi}_n(\omega))} (\phi_n - \phi) dF(\phi_n)$ reflects our assumption that the scrap value is random. It vanishes in a game of complete information such as Ericson & Pakes (1995), where scrap values are constant across firms and periods.

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

⁶If the support of $F(\cdot)$ is bounded, we define $F^{-1}(0)$ ($F^{-1}(1)$) to be the infimum (supremum) of the support.

is in state ω and the firm has drawn setup cost ϕ_n^e is

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

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.⁷ Depending on the entry decision $\tilde{\chi}_n^e(\omega, \phi_n^e)$, 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 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

⁷It 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).

(5) yields

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

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

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

$h_n(\cdot)$ is called firm n 's return (Denardo 1967, p. 166) or local income function (Whitt 1980, p. 35). 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 .

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

3 An Introductory Example

In this section, we first give an example in which a symmetric equilibrium in pure entry/exit strategies fails to exist in a setting with complete information such as Ericson & Pakes (1995).⁸ We then demonstrate that incorporating firm heterogeneity in the form of random scrap values/setup costs into the Ericson & Pakes (1995) model restores existence. We close this section with a brief discussion of computational issues.

We set $N = 2$ and $M = 1$. This implies that the industry is either a monopoly (states

⁸We defer a formal definition of our symmetry notion to Section 4.2.

(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.⁹ 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 $\pi(\omega_2, \omega_1)$. 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.”

Example: Deterministic scrap values/setup costs. 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 $\chi(1, 1) \in \{0, 1\}$, its value function is defined by the Bellman equation

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

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

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

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 $\chi(1, 1) = 1$

⁹A similar example can be constructed to demonstrate that there may not exist a symmetric equilibrium in pure entry strategies.

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.

Although there is no symmetric equilibrium in pure exit strategies there is one in mixed exit strategies given by

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

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 monopolistic industry and prescribing that both firms should exit from a duopolistic one.

To restore existence we introduced in Section 2 random scrap values in addition to the random setup costs suggested by Pakes & McGuire (1994).¹⁰ We now modify the above example to illustrate the 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

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

¹⁰Pakes & McGuire (1994) also suggest imposing an exogenous order on exit decisions as a means for addressing convergence problems (p. 570). Our formulation treats entry and exit on an equal footing, i.e., we introduce randomness into both decisions as suggested by Gowrisankaran (1995).

parameter	$\pi(1, 1)$	$\pi(1, 2)$	ϕ	β
value	0	1	15	$\frac{1}{1.05}$

Table 1: Parameter values.

ϵ	$V(\omega_1, \omega_1)$	$V(\omega_1, \omega_2)$	$\xi(\omega_1, \omega_1)$	$\xi(\omega_1, \omega_2)$
10	21.15967144	23.81754397	0.784835743	0.884168761
5	18.04492189	21	0.780374511	1
1	15.7308879	21	0.854919822	1
0.1	15.07621885	21	0.873034083	1
0.01	15.00765282	21	0.874803972	1
0.001	15.00076559	21	0.874980403	1
10^{-6}	15.00000077	21	0.87499998	1

Table 2: Equilibrium with random scrap values.

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 $\tilde{\xi}(\omega) = F\left(\frac{\bar{\phi}(\omega) - \phi}{\epsilon}\right)$ where $\bar{\phi}(\omega)$ is as in equations (11) and (12) except that $\chi(\omega)$ is replaced by $\xi(\omega)$.¹¹ 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 $0 < \xi(1, 1) < 1$ and $0 < \xi(1, 2) < 1$, and if $\epsilon \leq 5$, it involves $0 < \xi(1, 1) < 1$ and $\xi(1, 2) = 1$. Table 2 describes the equilibrium for various values of ϵ . Given the parameter values in Table 1, the symmetric equilibrium in mixed strategies of the game of complete information is $V(1, 2) = 21$, $V(1, 1) = 15$, $\xi(1, 2) = 1$, and $\xi(1, 1) = \frac{7}{8} = 0.875$. As Table 2 shows, the equilibrium with random scrap values converges to the equilibrium in mixed strategies as ϵ approaches zero. In the next section, we show that existence and convergence are general

¹¹To 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.

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. In comparison to the Ericson & Pakes (1995) model the Bellman equation now 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. Similarly, the cutoff scrap value/setup cost that determines the probability an incumbent firm remains in the industry/a potential entrant enters the industry is easily calculated. Thus introducing incomplete information into the Ericson & Pakes (1995) model adds essentially nothing to the computational burden.

In fact, the equilibrium of the game of incomplete information can be computed using a slightly modified version of the algorithm that Pakes & McGuire (1994) developed. This algorithm works iteratively. In the l th iteration, it 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} . In the context of the above example 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 $\lambda \in (0, 1)$. There is a clear trade-off between convergence and speed. 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.

The conclusion to be drawn from this discussion depends on the modeler's objective.

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.

If it is to calculate an equilibrium in mixed entry/exit strategies of the game of complete information, then the game of incomplete information may be a feasible way to obtain an approximation to such an equilibrium. The original algorithm must fail in the case in which only an equilibrium in (nondegenerate) mixed strategies exists, for that algorithm can neither exactly compute such an equilibrium nor can it closely approximate it. Slow convergence may therefore be the price payed to compute an approximation to an equilibrium in mixed strategies. If, on the other hand, random scrap values/setup costs are thought to be an accurate description of industry fundamentals, then the modified algorithm can be applied to compute an equilibrium in cutoff entry/exit strategies. In sharp contrast to the game of complete information, the search for an equilibrium is never in vain because an equilibrium in cutoff entry/exit strategies is guaranteed to exist. 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 in cutoff entry/exit and pure investment strategies. We first establish the existence of a possibly asymmetric and nonanonymous equilibrium. The proof extends Whitt (1980) to a setting with incomplete information. 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 several ways. We first show that a symmetric and anonymous equilibrium exists. This is essential from a computational viewpoint because symmetry and anonymity substantially reduce the computational burden. 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.

4.1 Existence

We begin by stating and discussing a series of assumptions. Our first assumption 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.¹²

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) 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 .*

¹²In models of learning-by-doing (Cabral & Riordan 1994, Benkard 2003), 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.

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 that entry and exit decisions are implemented before investment decisions are carried out. Hence, firm n chooses $x_n(\omega)$ in order to maximize $h_n(\omega, 1, x_n(\omega), u_{-n}(\omega), V_n)$ (see 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. We thus 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 a computationally tractable 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. 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.¹³

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

¹³Such an algorithm has indeed been suggested by Rust (1994).

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), and thus 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 $(H_uV_n)(\omega)$ 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 V_n . Like $H_{n,u_{-n}}^*$, $H_{n,u}$ is a contraction with modulus β that takes bounded vectors into bounded vectors. Hence, there exists a unique $V_{n,u} \in \mathcal{V}_n$ 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 equation (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 begin-

ning 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 .¹⁴ 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 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

¹⁴We set $\mathcal{W}_n = \mathcal{V}_n$ to obtain a special case of Whitt's (1980) lemma.

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: there exists a $u \in \mathcal{U}$ such that $u \in \Upsilon(u)$. ■

Assumption 3 is somewhat stronger than what is actually needed in the proof of proposition 1. Recall from equation (18) that $\Upsilon_n(\cdot)$ depends on $V_{n,u_{-n}}^*$, i.e., firm n chooses its best reply given that its future returns are described by the maximal return function. To ensure the existence of an equilibrium in pure investment strategies it therefore suffices that the local income function $h_n(\dots, 1, x_n(\omega), \dots)$ is maximized at a unique investment choice $x_n(\omega)$ for all possible maximal return functions $V_{n,u_{-n}}^*$ rather than for all possible value functions V_n . Hence, if \bar{V}^* and \underline{V}^* denote respectively a loose upper and a loose lower bound on the maximal return function $V_{n,u_{-n}}^*$, then assumption 3 can be weakened:

Assumption 4 *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 $u_{-n}(\omega)$, $V_n \in [\underline{V}^*, \bar{V}^*]^{|S|}$, ω , and all n .*

The bounds \bar{V}^* and \underline{V}^* are readily derived from the bounds $\bar{\phi}$, $\bar{\phi}^e$, and $\bar{\pi}$ on the fundamentals $F(\cdot)$, $F^e(\cdot)$, and $\pi_n(\cdot)$, respectively. The best possible net present value of the current and future cash flow that any firm, be it an incumbent or an entrant, can realize is no greater than

$$\bar{V}^* = \bar{\phi}^e + \frac{\bar{\pi}}{1 - \beta} + \bar{\phi}, \quad (19)$$

which is the sum of (i) a bound on its entry subsidy (i.e., negative setup cost), (ii) the capitalized value of remaining in the best possible state forever, and (iii) a bound on its scrap value. Because a firm always has the option of investing zero, it can guarantee that the net present value of its current and future cash flow is no worse than

$$\underline{V}^* = -\bar{\phi}^e - \frac{\bar{\pi}}{1 - \beta} - \bar{\phi}, \quad (20)$$

which is the sum of (i) a bound on its setup cost, (ii) the capitalized value of remaining in the worst possible state forever, and (iii) a bound on its exit tax (i.e., negative scrap value).

Replacing assumption 3 by assumption 4, we immediately obtain

Corollary 1 *Under assumptions 1, 2, and 4, an equilibrium exists in cutoff entry/exit and pure investment strategies.*

We use assumption 4 in the remainder of the paper. It turns out to be especially helpful in proving that UIC admissibility of the transition function $P(\cdot)$ guarantees existence of a computationally tractable equilibrium.

4.2 Symmetry and Anonymity

In Section 4.1 we established the existence of a possibly asymmetric and nonanonymous equilibrium. We now show that if an additional symmetry and anonymity assumption is satisfied, then a symmetric and anonymous equilibrium exists. In a symmetric equilibrium, if $V_1(\omega_1, \dots, \omega_N)$ denotes firm 1's value function, then firm n 's value function is given by $V_n(\omega_1, \dots, \omega_{n-1}, \omega_n, \omega_{n+1}, \dots, \omega_N) = V_1(\omega_n, \dots, \omega_{n-1}, \omega_1, \omega_{n+1}, \dots, \omega_N)$, and similarly for the policy function. Symmetry allows us to focus on the problem of firm 1. This problem can be further simplified by invoking exchangeability or anonymity. Anonymity says that firm 1 does not care about the identity of its competitors. Hence, $V_1(\omega_1, \omega_2, \dots, \omega_k, \dots, \omega_l, \dots, \omega_N) = V_1(\omega_1, \omega_2, \dots, \omega_l, \dots, \omega_k, \dots, \omega_N)$ for all $k \geq 2$ and $l \geq 2$, and similarly for the policy function. This, in effect, considerably reduces the size of the state space.

We begin by formalizing our symmetry and anonymity assumption.

Assumption 5 *The local income functions are symmetric and exchangeable, i.e.,*

$$\begin{aligned} & h_n(\omega_1, \dots, \omega_{n-1}, \omega_n, \omega_{n+1}, \dots, \omega_N, u_1(\omega), \dots, u_{n-1}(\omega), u_n(\omega), u_{n+1}(\omega), \dots, u_N(\omega), V_n) \\ &= h_1(\omega_n, \dots, \omega_{n-1}, \omega_1, \omega_{n+1}, \dots, \omega_N, u_n(\omega), \dots, u_{n-1}(\omega), u_1(\omega), u_{n+1}(\omega), \dots, u_N(\omega), V_1) \end{aligned}$$

for all symmetric value functions and all n , and

$$\begin{aligned} & h_1(\omega_1, \omega_2, \dots, \omega_k, \dots, \omega_l, \dots, \omega_N, u_1(\omega), u_2(\omega), \dots, u_k(\omega), \dots, u_l(\omega), \dots, u_N(\omega), V_1) \\ &= h_1(\omega_1, \omega_2, \dots, \omega_l, \dots, \omega_k, \dots, \omega_N, u_1(\omega), u_2(\omega), \dots, u_l(\omega), \dots, u_k(\omega), \dots, u_N(\omega), V_1) \end{aligned}$$

for all exchangeable value functions, $k \geq 2$, and all $l \geq 2$.

While we have stated assumption 5 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.

Lemma 1 *Assumption 5 is satisfied provided that (i) the profit functions are symmetric and anonymous, i.e.,*

$$\pi_n(\omega_1, \dots, \omega_{n-1}, \omega_n, \omega_{n+1}, \dots, \omega_N) = \pi_1(\omega_n, \dots, \omega_{n-1}, \omega_1, \omega_{n+1}, \dots, \omega_N)$$

for all n and

$$\pi_1(\omega_1, \omega_2, \dots, \omega_k, \dots, \omega_l, \dots, \omega_N) = \pi_1(\omega_1, \omega_2, \dots, \omega_l, \dots, \omega_k, \dots, \omega_N)$$

for all $k \geq 2$ and all $l \geq 2$; and (ii) the transition function is anonymous, i.e.,

$$\begin{aligned}
& P(\omega'_1, \dots, \omega'_k, \dots, \omega'_l, \dots, \omega'_N, \omega_1, \dots, \omega_k, \dots, \omega_l, \dots, \omega_N, \\
& \quad \chi_1(\omega, \phi_1), \dots, \chi_k(\omega, \phi_k), \dots, \chi_l(\omega, \phi_l), \dots, \chi_N(\omega, \phi_N), \\
& \quad x_1(\omega), \dots, x_k(\omega), \dots, x_l(\omega), \dots, x_N(\omega)) \\
= & P(\omega'_1, \dots, \omega'_l, \dots, \omega'_k, \dots, \omega'_N, \omega_1, \dots, \omega_l, \dots, \omega_k, \dots, \omega_N, \\
& \quad \chi_1(\omega, \phi_1), \dots, \chi_l(\omega, \phi_l), \dots, \chi_k(\omega, \phi_k), \dots, \chi_N(\omega, \phi_N), \\
& \quad x_1(\omega), \dots, x_l(\omega), \dots, x_k(\omega), \dots, x_N(\omega))
\end{aligned}$$

for all $k \geq 1$ and all $l \geq 1$.

The proof of lemma 1 is straightforward but tedious and therefore omitted. Note that in the special case of independent transitions, condition (ii) of lemma 1 is satisfied whenever the factors $P_n(\cdot)$ of the transition function $P(\cdot)$ are the same across firms, i.e., $P_n(\omega'_n, \omega_n, \chi_n(\omega, \phi_n), x_n(\omega)) = P_1(\omega'_n, \omega_n, \chi_n(\omega, \phi_n), x_n(\omega))$ for all n .

Together with assumptions 1, 2, and 4 in Section 4.1, assumption 5 ensures existence of a symmetric and anonymous equilibrium.

Proposition 2 *Under assumptions 1, 2, 4, and 5, a symmetric and anonymous 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. We enforce anonymity by redefining the state space to make it impossible for firm 1 to tailor its policy to the identity of its rivals. The argument in Section 4.1 implies that there exists a fixed point to a suitably defined best-reply correspondence of firm 1. The fixed point is used to construct a candidate equilibrium by specifying symmetric and anonymous policies for all firms. Since the associated value functions are also symmetric and anonymous, the argument is completed by exploiting the symmetry and anonymity of the local income functions to show that no firm has an incentive to deviate from the candidate equilibrium.

Before stating the proof we need to establish some notation. We begin with defining the reduced state space. Consider firm n and state ω . Define $\sigma_n = (\sigma_{n,1}, \dots, \sigma_{n,M}, \sigma_{n,M+1}) \in \Sigma \subset \{0, 1, \dots, N-1\}^{M+1}$, where $\sigma_{n,m}$ denotes the number of competitors of firm n that are in state m (excluding firm n), and 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 $\tau_n(\omega) = (\omega_n, \sigma_n)$. 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 the reduced state space is considerably smaller than the full state space. In fact, it has just $|S^\circ| = (M+1) \binom{M+N-1}{N-1} < (M+1)^N = |S|$ states.¹⁵

¹⁵Gowrisankaran (1999) develops an algorithm for the efficient representation of the reduced state space.

Define also the inverse function $\tau_n^{-1} : S^\circ \rightarrow S$ such that $\omega = \tau_n^{-1}(\omega_n, \sigma_n)$ is an arbitrary (but fixed) selection from the set $\{\omega \mid (\omega_n, \sigma_n) = \tau_n(\omega)\}$. For example, the state of the industry $\omega = \tau_n^{-1}(\omega_n, \sigma_n)$ may satisfy $\omega_1 \leq \omega_2 \leq \dots \leq \omega_{n-1} \leq \omega_{n+1} \leq \dots \leq \omega_N$. Note that, if $\tilde{\omega} = \tau_n^{-1}(\tau_n(\omega))$, then $\tilde{\omega}$ is obtained from ω by rearranging the elements of ω_{-n} . More specifically, $\tilde{\omega}_n = \omega_{\kappa_n}$ for some permutation $(\kappa_1, \dots, \kappa_N)$ of $(1, \dots, N)$ with $\kappa_n = n$, i.e., $\tilde{\omega}$ may differ from ω in all but the n th element.

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)). 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 by construction satisfies anonymity. 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

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

where Λ_n maps a value (or policy) function defined on the reduced state space to a value (or policy) function defined on the full state space such that $V_n = \Lambda_n(V_1^\circ)$ iff

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

for all ω .

The next step is to construct a candidate equilibrium. First, define the maximal return operator $H_{1, u_{-1}^\circ}^{\circ*} : \mathcal{V}_1^\circ \rightarrow \mathcal{V}_1^\circ$ pointwise by

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

where, to enforce symmetry, we take

$$u_{-1}^\circ = \underbrace{(u_1^\circ, \dots, u_1^\circ)}_{N-1 \text{ times}}$$

The maximal return function $V_{1, u_{-1}^\circ}^{\circ*}$ satisfies $V_{1, u_{-1}^\circ}^{\circ*} = H_{1, u_{-1}^\circ}^{\circ*} V_{1, u_{-1}^\circ}^{\circ*}$. Note that (i) $V_{1, u_{-1}^\circ}^{\circ*}$ is well-defined due to assumption 1 and (ii) that there is no circularity involved in its construction because u_1° (and hence u_{-1}°) is taken as given. Second, define the best-reply

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

$$\Upsilon_1^\circ(u_1^\circ) = \left\{ \tilde{u}_1^\circ \in \mathcal{U}_1^\circ : \tilde{u}_1^\circ(\omega_1, \sigma_1) \in \arg \sup_{\tilde{u}_1^\circ(\omega_1, \sigma_1) \in \mathcal{U}_1^\circ(\omega_1, \sigma_1)} h_1^\circ((\omega_1, \sigma_1), \tilde{u}_1^\circ(\omega_1, \sigma_1)), \right. \\ \left. u_1^\circ(\tau_2(\tau_1^{-1}(\omega_1, \sigma_1))), \dots, u_1^\circ(\tau_N(\tau_1^{-1}(\omega_1, \sigma_1))), V_{1, u_{-1}^\circ}^{\circ*} \text{ for all } (\omega_1, \sigma_1) \right\}. \quad (21)$$

where the maximal return function $V_{1, u_{-1}^\circ}^{\circ*}$ satisfies $V_{1, u_{-1}^\circ}^{\circ*} = H_{1, u_{-1}^\circ}^{\circ*} V_{1, u_{-1}^\circ}^{\circ*}$. Under assumptions 1, 2, and 4, 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. Third, construct a candidate equilibrium by using u_1° to define firm n 's policy function on the full state space to be

$$u_n = \Lambda_n(u_1^\circ). \quad (22)$$

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

The final step is to show that no firm has an incentive to deviate from the candidate equilibrium.

Proof of Proposition 2. Let $V_{1, u_{-1}^\circ}^{\circ*}$ denote the maximal return function corresponding to u_1° , i.e., $V_{1, u_{-1}^\circ}^{\circ*}$ satisfies $V_{1, u_{-1}^\circ}^{\circ*} = H_{1, u_{-1}^\circ}^{\circ*} V_{1, u_{-1}^\circ}^{\circ*}$ with $u_{-1}^\circ = (u_1^\circ, \dots, u_1^\circ)$. Define firm n 's value function on the full state space by

$$V_{n, u_{-n}}^* = \Lambda_n(V_{1, u_{-1}^\circ}^{\circ*}). \quad (23)$$

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

The key to the proof is to note that the policy functions in equation (22) satisfy a property stronger than symmetry and anonymity: if $\check{\omega}$ is obtained by rearranging the elements of ω , i.e., $\check{\omega}_n = \omega_{\kappa_n}$ for some permutation $(\kappa_1, \dots, \kappa_N)$ of $(1, \dots, N)$, then we have

$$\begin{aligned} u_n(\check{\omega}) &= u_1^\circ(\tau_n(\check{\omega})) \\ &= u_1^\circ(\tau_n(\check{\omega}_1, \dots, \check{\omega}_{n-1}, \check{\omega}_n, \check{\omega}_{n+1}, \dots, \check{\omega}_N)) \\ &= u_1^\circ(\tau_n(\omega_{\kappa_1}, \dots, \omega_{\kappa_{n-1}}, \omega_{\kappa_n}, \omega_{\kappa_{n+1}}, \dots, \omega_{\kappa_N})) \\ &= u_1^\circ(\tau_{\kappa_n}(\omega)) = u_{\kappa_n}(\omega). \end{aligned} \quad (24)$$

In particular, if $\kappa_n = n$, then we have $u_n(\check{\omega}) = u_n(\omega)$.

Equation (24) allows us to show that the problem of firm n in state ω is identical to the problem of firm 1 in state $\check{\omega}$ where $\check{\omega}$ is obtained by switching the first with the n th element of ω , i.e., $\check{\omega}_n = \omega_{\kappa_n}$ for $(\kappa_1, \kappa_2, \dots, \kappa_{n-1}, \kappa_n, \kappa_{n+1}, \dots, \kappa_N) = (n, 2, \dots, n-1, 1, n+1, \dots, N)$.

In fact, we have $\mathcal{U}_1(\tilde{\omega}) = \mathcal{U}_n(\omega)$ and

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

where the first equality follows from equation (24) and the second uses the symmetry of the value and local income functions. It thus remains to show the following:

$$\begin{aligned} u_1(\omega) \in \arg \sup_{\tilde{u}_1(\omega) \in \mathcal{U}_1(\omega)} & h_1(\omega, \tilde{u}_1(\omega), u_{-1}(\omega), V_{1,u_{-1}}^*), \\ V_{1,u_{-1}}^*(\omega) = \sup_{\tilde{u}_1(\omega) \in \mathcal{U}_1(\omega)} & h_1(\omega, \tilde{u}_1(\omega), u_{-1}(\omega), V_{1,u_{-1}}^*) \end{aligned}$$

for all ω . That is, (i) taking the value function in equation (23) as given, firm 1 has no incentive to deviate from the policy function in equation (22) and (ii) the value function in equation (23) coincides with the maximal return function of firm 1.

A state ω is called *canonical* iff $\omega = \tau_1^{-1}(\omega_1, \sigma_1)$ for some (ω_1, σ_1) . Fix a state ω . If state ω is canonical, then we have

$$\begin{aligned} & h_1(\omega, u_1(\omega), u_2(\omega), \dots, u_N(\omega), V_{1,u_{-1}}^*) \\ = & h_1(\omega, u_1^\circ(\tau_1(\omega)), u_1^\circ(\tau_2(\omega)), \dots, u_1^\circ(\tau_N(\omega)), \Lambda_1(V_{1,u_{-1}}^{\circ*})) \\ = & h_1(\tau_1^{-1}(\omega_1, \sigma_1), u_1^\circ(\tau_1(\tau_1^{-1}(\omega_1, \sigma_1))), u_1^\circ(\tau_2(\tau_1^{-1}(\omega_1, \sigma_1))), \dots, u_1^\circ(\tau_N(\tau_1^{-1}(\omega_1, \sigma_1))), \Lambda_1(V_{1,u_{-1}}^{\circ*})) \\ = & h_1(\tau_1^{-1}(\omega_1, \sigma_1), u_1^\circ(\omega_1, \sigma_1), u_1^\circ(\tau_2(\tau_1^{-1}(\omega_1, \sigma_1))), \dots, u_1^\circ(\tau_N(\tau_1^{-1}(\omega_1, \sigma_1))), \Lambda_1(V_{1,u_{-1}}^{\circ*})) \\ = & h_1^\circ((\omega_1, \sigma_1), u_1^\circ(\omega_1, \sigma_1), u_1^\circ(\tau_2(\tau_1^{-1}(\omega_1, \sigma_1))), \dots, u_1^\circ(\tau_N(\tau_1^{-1}(\omega_1, \sigma_1))), V_{1,u_{-1}}^{\circ*}), \end{aligned}$$

where the first equality follows from equations (22) and (23) and the last from the definition of the local income function $h_1^\circ(\cdot)$. Moreover, we have $\mathcal{U}_1(\omega) = \mathcal{U}_1(\tau_1^{-1}(\omega_1, \sigma_1)) = \mathcal{U}_1^\circ(\omega_1, \sigma_1)$. The fact that $u_1^\circ(\omega_1, \sigma_1) \in \mathcal{U}_1^\circ(\omega_1, \sigma_1)$ is optimal given $u_1^\circ(\tau_2(\tau_1^{-1}(\omega_1, \sigma_1))), \dots, u_1^\circ(\tau_N(\tau_1^{-1}(\omega_1, \sigma_1)))$ and $V_{1,u_{-1}}^{\circ*}$ (see the definition of $\Upsilon_1^\circ(\cdot)$ in equation (21)) therefore implies that $u_1(\omega) \in \mathcal{U}_1(\omega)$ is optimal given $u_2(\omega), \dots, u_N(\omega)$ and $V_{1,u_{-1}}^*$. Hence, (i) firm 1 has no incentive to deviate in the canonical state ω . (ii) now follows from (i) using $V_{1,u_{-1}}^*(\omega) = V_{1,u_{-1}}^*(\tau_1^{-1}(\omega_1, \sigma_1)) = V_{1,u_{-1}}^{\circ*}(\omega_1, \sigma_1)$.

If state ω is *not* canonical, then a canonical state $\tilde{\omega}$ can be obtained from ω by rearranging the elements of ω_{-1} . Formally, $\tilde{\omega}_n = \omega_{\kappa_n}$ for some permutation $(\kappa_1, \dots, \kappa_N)$ of $(1, \dots, N)$ with $\kappa_1 = 1$, i.e., $\tilde{\omega}$ may differ from ω in all but the first element. We have

$\mathcal{U}_1(\tilde{\omega}) = \mathcal{U}_1(\omega)$ and

$$\begin{aligned}
& h_1(\tilde{\omega}, u_1(\tilde{\omega}), u_2(\tilde{\omega}), \dots, u_N(\tilde{\omega}), V_{1,u_{-1}}^*) \\
&= h_1(\tilde{\omega}, u_1(\omega), u_{\kappa_2}(\omega), \dots, u_{\kappa_N}(\omega), V_{1,u_{-1}}^*) \\
&= h_1(\tilde{\omega}_1, \tilde{\omega}_2, \dots, \tilde{\omega}_N, u_1(\omega), u_{\kappa_2}(\omega), \dots, u_{\kappa_N}(\omega), V_{1,u_{-1}}^*) \\
&= h_1(\omega_1, \omega_{\kappa_2}, \dots, \omega_{\kappa_N}, u_1(\omega), u_{\kappa_2}(\omega), \dots, u_{\kappa_N}(\omega), V_{1,u_{-1}}^*) \\
&= h_1(\omega_1, \omega_2, \dots, \omega_N, u_1(\omega), u_2(\omega), \dots, u_N(\omega), V_{1,u_{-1}}^*),
\end{aligned}$$

where the first equality follows from equation (24) and the last equality uses the anonymity of the value and local income functions. That is, the problem of firm 1 in the non-canonical state ω is identical to the problem of firm 1 in the canonical state $\tilde{\omega}$. But we have already shown that (i) and (ii) hold in a canonical state. ■

4.3 Convergence

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

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

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

Corollary 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 3 *Suppose assumptions 1, 2, and 4 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 $\{V_{u^l}^{\epsilon^l}\}$ 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 $\{V_{u^l}^{\epsilon^l}\}$ 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 3 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 3 applies to the subsequential limit. This establishes

Corollary 2 *Under assumptions 1, 2, and 4, 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 4 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 $u_{-n}(\omega)$, $V_n \in [\underline{V}^*, \bar{V}^*]^{|S|}$, 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 4 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 4 that if this condition on the model's primitives is satisfied and if the upper bounds \bar{x} and \bar{x}^e on investment are sufficiently large, 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. We close this section by showing that UIC admissibility rules out equilibria that involve mixed investment strategies.

Condition 1 *The transition function $P(\cdot)$ is unique investment choice (UIC) admissible if, for all $\xi_{-n}(\omega)$, $x_n(\omega)$, $V_n \in [\underline{V}^*, \bar{V}^*]^{|S|}$, ω , and all n , 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 $h_n(\cdot)$ can be written in a separable form as*

$$\sum_{\omega'_n \in \Omega} \left[K_n(\omega'_n, \omega, u_{-n}(\omega), V_n) Q_n(\omega'_n, \omega, x_n(\omega)) + L_n(\omega'_n, \omega, u_{-n}(\omega), V_n) \right], \quad (25)$$

where, for all ω'_n , $Q_n(\omega'_n, \omega, x_n(\omega))$ is continuously differentiable and

$$\frac{d}{dx_n(\omega)} Q_n(\omega'_n, \omega, x_n(\omega)) = \frac{d_{n,\omega,\omega'_n}}{a_{n,\omega} x_n(\omega)^2 + b_{n,\omega} x_n(\omega) + c_{n,\omega}} \quad (26)$$

with either (i) $a_{n,\omega} > 0$ and $b_{n,\omega} \geq 0$ or (ii) $a_{n,\omega} = 0$ and $b_{n,\omega} \neq 0$.

UIC admissibility ensures that firm n 's local income function $h_n(\dots, 1, x_n(\omega), \dots)$ has at most one extreme point in the interval $(0, \bar{x})$ (or in the interval $(0, \bar{x}^e)$ if firm n is an entrant rather than an incumbent). Taken by itself, UIC admissibility does not suffice to ensure a unique investment choice because investing zero may be just as good as investing \bar{x} (or \bar{x}^e). To rule this out, we stipulate that the upper bounds \bar{x} and \bar{x}^e on investment are large enough never to constrain firms' optimal investment choices. Specifically, we assume that \bar{x} and \bar{x}^e are larger than $\beta(\bar{V}^* - \underline{V}^*)$ where \bar{V}^* and \underline{V}^* are the bounds on the maximal return function $V_{n,u_{-n}}^*$ given in equations (19) and (20). This is mathematically innocuous since \bar{x} and \bar{x}^e can always be chosen to be arbitrarily large. It is, however, a genuine constraint to the extent that it rules out models in which investment choices are limited, for example

because of cash constraints.¹⁶

We are now ready to state our main result establishing that a computationally tractable equilibrium exists in the modified Ericson & Pakes (1995) model.

Proposition 4 *Suppose assumptions 1 and 2 hold. If the transition function $P(\cdot)$ is UIC admissible and \bar{x} and \bar{x}^e are finite and larger than $\beta(\bar{V}^* - \underline{V}^*)$, then an equilibrium exists in cutoff entry/exit and pure investment strategies. If in addition assumption 5 holds, then a symmetric and anonymous equilibrium exists in cutoff entry/exit and pure investment strategies.*

Proof. In light of corollary 1 and proposition 2, it suffices to show that assumption 4 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.

We begin by ruling out \bar{x} as an optimal investment choice. To do so, we show that it is always better to invest zero than to invest \bar{x} . By construction of \bar{V}^* and \underline{V}^* ,

$$h_n(\dots, 1, \bar{x}, \dots) \leq \pi_n(\omega) - \bar{x} + \beta\bar{V}^*, \quad h_n(\dots, 1, 0, \dots) \geq \pi_n(\omega) - 0 + \beta\underline{V}^*.$$

Hence, $h_n(\dots, 1, \bar{x}, \dots) - h_n(\dots, 1, 0, \dots) \leq -\bar{x} + \beta(\bar{V}^* - \underline{V}^*) < 0$ where the last inequality follows from our assumption that $\bar{x} > \beta(\bar{V}^* - \underline{V}^*)$. This implies first that \bar{x} cannot be an optimal investment choice and second that $h_n(\dots, 1, x_n(\omega), \dots)$ must be decreasing somewhere on $[0, \bar{x}]$.

Next we differentiate $h_n(\dots, 1, x_n(\omega), \dots)$ with respect to $x_n(\omega)$. Since $P(\cdot)$ is UIC admissible, the FOC for an unconstrained solution to firm n 's investment problem is

$$0 = -1 + \beta \sum_{\omega'_n \in \Omega} K_n(\omega'_n, \omega, u_{-n}(\omega), V_n) \frac{d_{n,\omega,\omega'_n}}{a_{n,\omega}x_n(\omega)^2 + b_{n,\omega}x_n(\omega) + c_{n,\omega}}.$$

Simplifying yields the quadratic equation $0 = a_{n,\omega}x_n(\omega)^2 + b_{n,\omega}x_n(\omega) + e_{n,\omega,\omega'_n}$ where

$$e_{n,\omega,\omega'_n} = \left\{ c_{n,\omega} - \beta \sum_{\omega'_n \in \Omega} K_n(\omega'_n, \omega, u_{-n}(\omega), V_n) d_{n,\omega,\omega'_n} \right\}.$$

Case (i): $a_{n,\omega} > 0$ and $b_{n,\omega} \geq 0$. Suppose first that $b_{n,\omega}^2 - 4a_{n,\omega}e_{n,\omega,\omega'_n} \geq 0$. Then there are two roots (which may coincide) to the quadratic equation: the smaller root is negative because $b_{n,\omega} \geq 0$ whereas the larger root,

$$\hat{x}_n(\omega) = \frac{-b_{n,\omega} + \sqrt{b_{n,\omega}^2 - 4a_{n,\omega}e_{n,\omega,\omega'_n}}}{2a_{n,\omega}},$$

¹⁶Of 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).

may be negative, zero, or positive.

The larger root $\hat{x}_n(\omega)$ is therefore the only candidate for an interior solution within $(0, \bar{x})$ to firm n 's investment problem. Moreover, the derivative of $h_n(\dots, 1, x_n(\omega), \dots)$ with respect to $x_n(\omega)$ can change sign at most once on $[0, \bar{x}]$: If it never changes sign, then $h_n(\dots, 1, x_n(\omega), \dots)$ must be strictly decreasing on $[0, \bar{x}]$. If it changes sign once on $[0, \bar{x}]$, then $\hat{x}_n(\omega) \in (0, \bar{x})$ must be either a local minimum or a local maximum.

1. $h_n(\dots, 1, x_n(\omega), \dots)$ may be strictly decreasing on the interval $[0, \bar{x}]$. Then the unique maximizer is 0.
2. $\hat{x}_n(\omega) \in (0, \bar{x})$ may be a local minimum. Then, on the interval $[0, \bar{x}]$, $h_n(\dots, 1, x_n(\omega), \dots)$ is strictly decreasing (increasing) to the left (right) of $\hat{x}_n(\omega)$. Because \bar{x} cannot be an optimal investment choice, the unique maximizer is 0.
3. $\hat{x}_n(\omega) \in (0, \bar{x})$ may be a local maximum. Then, on the interval $[0, \bar{x}]$, $h_n(\dots, 1, x_n(\omega), \dots)$ is strictly increasing (decreasing) to the left (right) of $\hat{x}_n(\omega)$. Hence, the unique maximizer is $\hat{x}_n(\omega)$.

Suppose next that $b_{n,\omega}^2 - 4a_{n,\omega}e_{n,\omega}\omega'_n < 0$, so that the quadratic equation has no real roots. Then $h_n(\dots, 1, x_n(\omega), \dots)$ must be either strictly decreasing or strictly increasing on $[0, \bar{x}]$. But we have already established that $h_n(\dots, 1, x_n(\omega), \dots)$ must be decreasing somewhere on $[0, \bar{x}]$; therefore it cannot possibly be strictly increasing on $[0, \bar{x}]$. Hence, $h_n(\dots, 1, x_n(\omega), \dots)$ is strictly decreasing on $[0, \bar{x}]$, and the unique maximizer is 0.

Case (ii): $a_{n,\omega} = 0$ and $b_{n,\omega} \neq 0$. The only candidate for an interior solution to firm n 's investment problem is

$$\hat{x}_n(\omega) = -\frac{e_{n,\omega}\omega'_n}{b_{n,\omega}}.$$

Again the derivative of $h_n(\dots, 1, x_n(\omega), \dots)$ with respect to $x_n(\omega)$ can change sign at most once on $[0, \bar{x}]$. The remainder of the proof is identical to case (i). ■

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 2002) where a firm's state describes its capacity. In each period, the firm decides how much to spend on an investment project in order to add to its capacity. If firm n invests $x_n(\omega) \geq 0$, then the probability that its investment project succeeds is

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

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

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

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 4.¹⁷

Without loss of generality, consider firm 1. For the case of $\omega'_1 = \omega_1$, the expected value of its future cash flow is

$$\begin{aligned}
& [\delta p_1 + (1 - \delta)(1 - p_1)]\{(1 - \delta)p_2 V_1(\omega_1, \omega_2 + 1) + [\delta p_2 + (1 - \delta)(1 - p_2)]V_1(\omega_1, \omega_2) \\
& \quad + \delta p_2 V_1(\omega_1, \omega_2 - 1)\} \\
= & \underbrace{\{(1 - \delta)p_2 V_1(\omega_1, \omega_2 + 1) + [\delta p_2 + (1 - \delta)(1 - p_2)]V_1(\omega_1, \omega_2) + \delta p_2 V_1(\omega_1, \omega_2 - 1)\}}_{K_1(\omega_1, \omega, x_2(\omega), V_1)} (2\delta - 1) \underbrace{p_1}_{Q_1(\omega_1, \omega, x_1(\omega))} \\
& + \underbrace{\{(1 - \delta)p_2 V_1(\omega_1, \omega_2 + 1) + [\delta p_2 + (1 - \delta)(1 - p_2)]V_1(\omega_1, \omega_2) + \delta p_2 V_1(\omega_1, \omega_2 - 1)\}}_{L_1(\omega_1, \omega, x_2(\omega), V_1)} (1 - \delta).
\end{aligned}$$

This expression satisfies the separability condition (25), as do the corresponding expressions for the cases of $\omega'_1 = \omega_1 + 1$ and $\omega'_1 = \omega_1 - 1$. In addition, the derivative condition (26) is satisfied because

$$\frac{d}{dx_1(\omega)} Q_1(\omega'_1, \omega, x_1(\omega)) = \frac{\alpha}{\alpha^2 x_1(\omega)^2 + 2\alpha x_1(\omega) + 1}.$$

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.

For the sake of brevity, we just spell out the expected value of firm 1's future cash flow

¹⁷Edge 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.

in case of $\omega'_1 = \omega_1$,

$$\begin{aligned}
& (1 - p_1) \{ (1 - \delta) p_2 V_1(\omega_1, \omega_2 + 1) + (1 - \delta) (1 - p_2) V_1(\omega_1, \omega_2) \} \\
& + p_1 \{ \delta p_2 V_1(\omega_1, \omega_2) + \delta (1 - p_2) V_1(\omega_1, \omega_2 - 1) \} \\
= & \underbrace{\{ - (1 - \delta) p_2 V_1(\omega_1, \omega_2 + 1) + [- (1 - \delta) (1 - p_2) + \delta p_2] V_1(\omega_1, \omega_2) + \delta (1 - p_2) V_1(\omega_1, \omega_2 - 1) \}}_{K_1(\omega_1, \omega, x_2(\omega), V_1)} \\
& \times \underbrace{p_1}_{Q_1(\omega_1, \omega, x_1(\omega))} + \underbrace{\{ (1 - \delta) p_2 V_1(\omega_1, \omega_2 + 1) + (1 - \delta) (1 - p_2) V_1(\omega_1, \omega_2) \}}_{L_1(\omega_1, \omega, x_2(\omega), V_1)},
\end{aligned}$$

and note that conditions (25) and (26) 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 $\eta_{n, \omega_n, \hat{\omega}_n}$.¹⁸ 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 (26). 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 $\hat{\omega}$.

¹⁸The 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 $\sum_{\hat{\omega}_n=1}^M \eta_{n, \omega_n, \hat{\omega}_n} = 0$.

Mixed investment strategies. From the outset we have restricted attention to pure investment strategies. Under assumption 4 the unique best reply to pure investment strategies $x_{-n}(\omega)$ is a pure investment strategy $x_n(\omega)$. An equilibrium in pure investment strategies therefore remains an equilibrium even if mixed investment strategies are allowed.

But we can say more: UIC admissibility rules out equilibria that involve mixed investment strategies. Consider the decision problem of firm n when the other firms mix on their actions $u_{-n}(\omega)$, including their investments, according to the distribution $\Psi(u_{-n}(\omega))$. Given that $P(\cdot)$ is UIC admissible, the expected value of firm n 's future cash flow can be written as

$$\begin{aligned}
& \int_{u_{-n}(\omega) \in \mathcal{U}_{-n}(\omega)} \sum_{\omega'_n \in \Omega} \left[\begin{array}{c} K_n(\omega'_n, \omega, u_{-n}(\omega), V_n) Q_n(\omega'_n, \omega, x_n(\omega)) \\ + L_n(\omega'_n, \omega, u_{-n}(\omega), V_n) \end{array} \right] d\Psi(u_{-n}(\omega)) \\
= & \sum_{\omega'_n \in \Omega} \left[\int_{u_{-n}(\omega) \in \mathcal{U}_{-n}(\omega)} K_n(\omega'_n, \omega, u_{-n}(\omega), V_n) d\Psi(u_{-n}(\omega)) Q_n(\omega'_n, \omega, x_n(\omega)) \right. \\
& \left. + \int_{u_{-n}(\omega) \in \mathcal{U}_{-n}(\omega)} L_n(\omega'_n, \omega, u_{-n}(\omega), V_n) d\Psi(u_{-n}(\omega)) \right] \\
= & \sum_{\omega'_n \in \Omega} \left[K_n^\diamond(\omega'_n, \omega, u_{-n}(\omega), V_n) Q_n(\omega'_n, \omega, x_n(\omega)) + L_n^\diamond(\omega'_n, \omega, u_{-n}(\omega), V_n) \right].
\end{aligned}$$

Hence, assumption 4 holds, and firm n 's best reply is unique and cannot be mixed.

6 Multiplicity

In this section we discuss three examples that show that there need not be a unique equilibrium that is symmetric and anonymous.

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.¹⁹ Products are undifferentiated and firms compete in prices subject to capacity 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 parameters values.

Figure 1 illustrates the value and policy functions of two equilibria. 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

¹⁹Because the transition function $P(\cdot)$ is UIC admissible, it is guaranteed that nonuniqueness is not due to a violation of Assumption 4.

parameter	M	\bar{q}_1	\bar{q}_2	\dots	\bar{q}_{10}	v	m	α	δ	β
value	10	0	5	\dots	45	1	10	2.375	0.03	$\frac{1}{1.05}$

Table 6: Parameter values.




Figure 1: Two equilibria.

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). Nevertheless, the two equilibria are qualitatively similar. This is especially troublesome because there is little hope of using empirical evidence to distinguish between these equilibria. However, the importance of multiplicity would greatly diminish if it could be shown that all equilibria are generally alike.

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 noting because many applications of Ericson & Pakes's (1995) framework have searched for multiple equilibria by selecting a single algorithm and varying the starting values. Our example makes clear that this approach may fail and 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. In contrast, Pakes &

McGuire (1994) have conjectured that nonuniqueness may result from firms' exit decisions. This is easily seen by slightly extending 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 ($\delta = 0$)). Then one symmetric equilibrium has both firms play the cutoff exit strategies from Section 3 in states $(1, 1)$, $(1, 2)$, $(2, 1)$, and $(2, 2)$. But there are other equilibria: For example, both firms play the cutoff exit strategies in states $(1, 1)$ and $(2, 2)$. In state $(1, 2)$ firm 1 exits for sure and firm 2 stays for sure whereas in state $(2, 1)$ firm 1 stays for sure and firm 2 exits for sure. This clearly shows how the symmetry requirement may fail to rule out all but one equilibrium.

Example: Product market competition. We close this section by noting that we treat the current 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 establishes a solid foundation of the Ericson & Pakes (1995) model of dynamic competition in an oligopolistic industry with investment, entry, and exit. We show that existence of a MPE in the Ericson & Pakes (1995) game of complete information requires mixed entry/exit strategies, contrary to their assertion. 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 a MPE in cutoff entry/exit strategies and is computationally no more demanding than the original game of complete information.

Adding random scrap values/setup costs formally leads to a dynamic stochastic game with compact and convex action spaces 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. Since computing mixed strategies over continuous actions is well beyond present computational capabilities, it is vital to ensure existence of a MPE 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 and anonymous MPE exists provided the model's primitives are symmetric and anonymous. This is a major result for two reasons. First, from a computational viewpoint, symmetry and anonymity are 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 and anonymous MPE an especially compelling solution concept because, in such a MPE, 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 and anonymous MPE in a broad class of dynamic stochastic games.²⁰ Our arguments are readily extended to arbitrary dynamic stochastic games.

Next we show that, as the distribution of the random scrap values/setup costs becomes degenerate, MPEs in cutoff entry/exit strategies converge to MPEs in mixed entry/exit strategies of the game of complete information. We have been unable to determine whether or not the approachability part of Harsanyi's (1973) purification theorem carries over from static games to dynamic stochastic games. That is, are all MPEs of the original game approached by some MPE of the perturbed game as the perturbation vanishes? We leave this as an open question for future research.

Finally, we provide the first example of multiple symmetric and anonymous MPEs in the literature spawned 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 MPE in dynamic stochastic games. Haller & Lagunoff (2000) show that the number of MPEs is generically finite and Amir (2002) shows that there exists a unique MPE in pure strategies in *finite* horizon games that satisfy the same monotonicity, supermodularity, and dominant-diagonal conditions that Curtat (1996) adopted in his earlier paper. More research along these lines is clearly needed.

References

- Amir, R. (2002), 'Complementarities and diagonal dominance in discounted stochastic games', *Annals of Operations Research* **114**(1), 39–56.
- Benkard, L. (2003), 'A dynamic analysis of the market for wide-bodied commercial aircraft', *Review of Economic Studies* **forthcoming**.
- Bergin, J. & Bernhardt, D. (1995), 'Anonymous sequential games: Existence and characterization of equilibria', *Economic Theory* **5**, 461–489.

²⁰Dutta & Sundaram (1992) show that there exists a symmetric MPE in two-player resource extraction games with a one-dimensional state space.

- Besanko, D. & Doraszelski, U. (2002), ‘Capacity dynamics and endogenous asymmetries in firm size’, *Rand Journal of Economics* **forthcoming**.
- Cabral, L. & Riordan, M. (1994), ‘The learning curve, market dominance, and predatory pricing’, *Econometrica* **62**(2), 1115–1140.
- Chakrabarti, S. (2003), ‘Pure strategy Markov equilibrium in stochastic games with a continuum of players’, *Journal of Mathematical Economics* **39**, 693–724.
- 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. & Markovich, S. (2003), Advertising dynamics and competitive advantage, Working paper, Hoover Institution, Stanford.
- Dunne, T., Roberts, M. & Samuelson, L. (1988), ‘Patterns of firm entry and exit in U.S. manufacturing’, *Rand Journal of Economics* **19**(4), 495–515.
- 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.
- Ericson, R. & Pakes, A. (1995), ‘Markov-perfect industry dynamics: A framework for empirical work’, *Review of Economic Studies* **62**, 53–82.
- Federgruen, A. (1978), ‘On N -person stochastic games with denumerable state space’, *Advances in Applied Probability* **10**, 452–471.
- Ghemawat, P. (1997), *Games businesses play*, MIT Press, Cambridge.
- Gowrisankaran, G. (1995), A dynamic analysis of mergers, Ph.D. thesis, Yale University, New Haven.
- Gowrisankaran, G. (1999), ‘Efficient representation of state spaces for some dynamic models’, *Journal of Economic Dynamics and Control* **23**, 1077–1098.
- Haller, H. & Lagunoff, R. (2000), ‘Genericity and Markovian behavior in stochastic games’, *Econometrica* **68**(5), 1231–1248.
- Harsanyi, J. (1973), ‘Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points’, *International Journal of Game Theory* **2**, 1–23.
- Judd, K. (1998), *Numerical methods in economics*, MIT Press, Cambridge.
- Langohr, P. (2003), Competitive convergence and divergence: Capability and position dynamics, Working paper, Northwestern University, Evanston.
- 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* **59**(5), 1261–1281.
- 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.
- 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.