Recursivity and the Estimation of Dynamic Games with Continuous Controls

João Granja
Giuseppe Forte

May 22, 2026

Abstract

We study the estimation of dynamic games with continuous control variables, such as investments in R&D, quality, and capacity. We use the recursive characterization of Markov Perfect Equilibria (MPE) to develop estimators that exploit the structure of optimal policies. In particular, we derive a pseudo maximum likelihood estimator for models with shocks to firms' marginal costs of investment. We evaluate the performance of these estimators in two Monte Carlo exercises, including a version of the Hashmi and van Biesebroeck (2016) model of innovation in the automobile industry extended to allow for entry and exit. We find that estimators based on recursive equilibrium conditions perform well and outperform the inequality estimator of Bajari, Benkard, and Levin (2007).

First draft: December 4, 2022.
Granja: University College London, joao.granja@ucl.ac.uk.
Forte: University College London and CeMMAP, uctpgfo@ucl.ac.uk.
We thank Dennis Kristensen, Robert Miller, Lars Nesheim, and Karl Schurter for helpful comments and suggestions.

1 Introduction

Many questions in Industrial Organization involve firm choices that have persistent effects on market structure. Such choices include, for example, investments in research and development, the choice of productive capacity, and the choice of product characteristics. These choices are inherently dynamic, are often taken in industries with few firms, and are naturally modeled as continuous variables. Therefore, their study necessitates dynamic oligopoly models with continuous controls.

This paper develops estimators for such models. The estimators we propose share the two-step structure common to many estimators of dynamic games with either discrete or continuous controls (Bajari et al., 2007; Aguirregabiria & Mira, 2007; Pakes, Ostrovsky, & Berry, 2007; Pesendorfer & Schmidt-Dengler, 2008), whereby the first step consists of estimating policy functions and state transitions. We depart from most of the prior literature in the second step, where we use the first-step estimates to construct an empirical analog of firms' Bellman equations. For models with a shock to firms' marginal costs of setting the continuous control, such as the general model of Bajari et al. (2007) and Hashmi and van Biesebroeck (2016), we exploit the empirical Bellman equations to derive a pseudo maximum likelihood estimator. For models without such shocks, such as the quality-ladder example of Bajari et al. (2007) and Ryan (2012), we use the empirical Bellman equations to solve for optimal behavior and form a nonlinear least squares estimator. Our approach, based on firms' Bellman equations and associated first-order conditions, stands in contrast to the estimator proposed by Bajari et al. (2007), which is based on value function inequalities stemming from the requirement that the observed policy's value be no less than that of any alternative. Because the estimators we propose make fuller use of the model structure, we expect them to exhibit improved econometric performance. We test that conjecture in two Monte Carlo designs.

The first design closely follows the quality ladder game in Bajari et al. (2007). In that model, firms invest to improve the quality of their products, conditioning their actions on their own and their rivals' quality levels, subject to a linear investment cost. Firms also make entry and exit decisions after observing private information entry costs and scrap values, respectively. The second design extends the Hashmi and van Biesebroeck (2016) model of innovation to allow for firm entry and exit. The environment is similar to the first design, but, importantly, firms' marginal costs of investment are subject to private information shocks. Importantly, this framework underlies an empirical study of the automobile industry, so the design serves to illustrate the performance of our estimators in a setting representative of models used in empirical applications. We compare our estimators against three implementations of the Bajari et al. (2007) estimator (henceforth, BBL), each using a different form of policy deviation, constructed by perturbing estimated policies: additive perturbations, as in Bajari et al. (2007) and Ryan (2012); multiplicative perturbations, as in Hashmi and van Biesebroeck (2016) and recommended by Srisuma (2013); and what we term asymptotic deviations, constructed from the asymptotic distribution of the empirical policies.

We find that our estimators perform substantially better than all three BBL variants in both designs. In the BBL design, our estimator recovers the investment cost coefficient essentially without bias, while the additive and multiplicative BBL variants underestimate it by 37 and 46 percent. Our estimates exhibit smaller finite-sample bias on the scrap value distribution and comparable bias on the entry cost distribution, and on both are 2–3 times more precise than the best BBL implementation. The asymptotic variant fails qualitatively, yielding investment cost estimates of the wrong sign and entry cost estimates orders of magnitude from the truth. In the HvB design, our estimator recovers the investment cost parameters with small bias, while all three BBL variants fail — returning estimates at the boundary, with the wrong sign, or off by an order of magnitude. These failures reflect a near-flat BBL objective in the investment cost parameters around the true parameter vector. On the scrap value scale parameter, our estimator is comparable to the best BBL implementation; on the entry cost scale, BBL variants overestimate by 30–70 percent, while our finite-sample bias is roughly 5 percent and our estimate is about twice as precise as the best BBL implementation. Notably, the asymptotic variant — which failed qualitatively in the BBL design — is the least-biased of the three BBL implementations here. Two main takeaways arise from these results. First, the performance of the BBL estimator depends markedly on the choice of deviations, and which form performs best can depend on the environment. Second, estimators based on recursive equilibrium conditions can markedly outperform the BBL estimator in empirically relevant settings.

The two papers most closely related to ours in methodological approach are Srisuma (2013) and Wang and Zhai (2025). Srisuma (2013) observes that the BBL inequalities may fail to identify structural parameters and proposes an estimator that makes use of agents' optimization problems in a two-step procedure.

Srisuma's estimator minimizes a distance between the observed and model-implied conditional action distributions — an expensive objective to compute, as evidenced by the simple static Monte Carlo designs in Srisuma (2013). Wang and Zhai (2025) independently develop a pseudo maximum likelihood estimator to study institutional landlords' acquisitions in the single-family housing market. Their framework accommodates multiple continuous controls, while our derivation makes explicit use of the monotonicity of firms' policy functions, a property for which we provide sufficient conditions.

This paper also relates to a rich literature estimating dynamic games with continuous controls. A number of empirical papers apply the BBL estimator, including Ryan (2012), Hashmi and van Biesebroeck (2016), Fowlie, Reguant, and Ryan (2016), Liu and Siebert (2022), and Egan, Hortaçsu, Kaplan, Sunderam, and Yao (2025). Others use recursive equilibrium conditions in estimation, similar in spirit to what we do. Jofre-Bonet and Pesendorfer (2003) use Bellman first-order conditions to identify the distribution of bidders' costs in a dynamic auction model. We instead use the empirical Bellman equations to estimate structural parameters governing firms' investment costs, entry costs, and scrap values — objects absent from their setting — while taking the cost shock distribution as known, as is standard in the literature (Bajari et al., 2007; Srisuma, 2013; Hashmi & van Biesebroeck, 2016). Lim and Yurukoglu (2018) estimate a dynamic game between a regulated utility controlling its capital level and the regulator by matching first-stage policies to those implied by an empirical Bellman equation. This is related to our estimation of the Bajari et al. (2007) model. Our work extends estimators based on recursive equilibrium conditions to models with private information shocks to firms' investment costs, and benchmarks their performance against the BBL estimator.

Our contributions are threefold. First, we derive estimators that exploit recursive equilibrium conditions for dynamic games with continuous controls. In particular, we derive a pseudo MLE estimator for models with private information shocks to firms' investment costs. That derivation rests on the monotonicity of firms' policy functions, a property for which we provide two sets of sufficient conditions. Second, our Monte Carlo exercises provide evidence that estimators based on recursive equilibrium conditions can substantially outperform the BBL estimator in designs representative of models used in empirical applications. Moreover, our results show that the performance of the BBL estimator can depend meaningfully on the choice of policy deviations. Finally, we provide an existence result for Markov Perfect Equilibria in dynamic games with continuous controls and private information shocks to firms' marginal costs of setting the continuous control, extending a result of Doraszelski and Satterthwaite (2010).

The rest of the paper is organized as follows. In Section 2 we discuss a general model of dynamic competition in an oligopolistic industry. In Section 3 we introduce the pseudo-MLE estimator based on firms' recursive optimality conditions and briefly review the BBL inequality estimator. In Section 4 we report the results from our Monte Carlo exercises. Section 5 concludes.

2 The Economic Model

We model the dynamic interaction between oligopolistic competitors. There are $\bar{N}$ firms in the market, including N incumbents and $\bar{N} - N$ potential entrants. $\bar{N}$ is a parameter of the model whereas N is an endogenous variable. Each firm has a characteristic $\xi_i \in \Xi$, where $\Xi \subseteq \mathbb{R} \cup \{-\infty\}$ contains $-\infty$ to represent an inactive firm. Time is discrete and the horizon is infinite. The state of the industry at time t is $\boldsymbol{\xi}_t = (\xi_{1t}, \ldots, \xi_{\bar{N}t})$.[1]

At the beginning of the period firm i earns flow profit $\pi_i(\boldsymbol{\xi}_t)$. If $\xi_i = -\infty$, then $\pi_i(\boldsymbol{\xi}_t) = 0$. Flow profits are typically modeled as the outcome of competition in static variables such as prices or quantities. We do not need to specify the underlying model that generates π_i, but rather treat these functions as parameters of the dynamic game. We assume that the functions π_i are symmetric, i.e., that

$$\pi_i(\xi_i, \xi_2, \ldots, \xi_{i-1}, \xi_1, \xi_{i+1}, \ldots, \xi_{\bar{N}}) = \pi_1(\boldsymbol{\xi}) =: \pi(\boldsymbol{\xi}) \quad \text{for all } i = 2, \ldots, \bar{N} \tag{1}$$

and

$$\pi(\xi_1, \boldsymbol{\xi}_{-1}) = \pi(\xi_1, \boldsymbol{\xi}_{p(-1)}) \tag{2}$$

for any permutation $p(-1)$ of the indices $2, \ldots, \bar{N}$ — see, e.g., Doraszelski and Satterthwaite (2010).[2]

After firms earn profits, incumbents privately observe scrap values $\rho_{it} \in \mathbb{R}_+$ and potential entrants privately observe entry costs $\kappa_{it} \in \mathbb{R}_+$. Scrap values and entry costs are independent and identically distributed (iid) across firms

[1]It is straightforward to accommodate exogenous states that capture, e.g., changing demand and/or cost conditions.

[2]These conditions are sometimes called, respectively, symmetry and anonymity — see, e.g., Doraszelski and Pakes (2007). Doraszelski and Satterthwaite (2010) call a set of functions symmetric if they satisfy both conditions. We adopt their terminology.

and periods, with distributions F_ρ and F_κ, respectively. Upon observing these random variables, firms simultaneously decide whether or not to be active in period $t+1$. We denote the decision to be active by $\alpha_{it} = 1$; choosing not to be active is represented by $\alpha_{it} = 0$. Firms who decide not to be active in the next period perish and are replaced by new potential entrants.

Besides entry and exit decisions, firms also invest to affect the evolution of their characteristics ξ_{it}. Investment choices are denoted by $x_{it} \in [0, \bar{x}]$, where $\bar{x} < \infty$ is an exogenously imposed upper bound on investment.[3] After firms make their entry and exit decisions, all firms that chose to be active in $t+1$ observe investment cost shocks $\nu_{it} \in \mathcal{S} \subseteq \mathbb{R}$, drawn iid across firms and periods from F_ν. They then simultaneously choose their levels of investment and incur costs $c(x_{it}, \nu_{it})$. We let $a_{it} = (\alpha_{it}, x_{it})$ denote the vector of firm i's actions in period t, which satisfies the restriction that $x_{it} = 0$ whenever $\alpha_{it} = 0$.[4]

Assumption 1 (Investment Cost (IC))**.** The investment cost function $c(x, \nu)$ is twice continuously differentiable, strictly increasing and convex in x, and the marginal cost of investment is strictly increasing in the cost shock ν, i.e.

$$\frac{\partial c(x,\nu)}{\partial x} > 0, \quad \frac{\partial^2 c(x,\nu)}{\partial x^2} \geq 0, \quad \text{and} \quad \frac{\partial^2 c(x,\nu)}{\partial \nu \partial x} > 0 \ .$$

Assumption 2 (Conditional Independence and No Spillovers (CINS))**.** The distribution of $\boldsymbol{\xi}_{t+1}$ conditional on $\boldsymbol{\xi}_t$ and firms' action profile $\boldsymbol{a}_t = (a_{1t}, \ldots, a_{\bar{N}t})$, where $a_{it} = (\alpha_{it}, x_{it})$, satisfies

$$F_{\boldsymbol{\xi}}(\boldsymbol{\xi}_{t+1} \mid \boldsymbol{\xi}_t, \boldsymbol{a}_t) = \prod_{i=1}^{\bar{N}} F_\xi(\xi_{i,t+1} \mid \xi_{i,t}, a_{it}) \ . \tag{3}$$

We restrict the transition function F_ξ so that firms move to state $\xi = -\infty$ only if they choose not to be active in the following period: for all $\xi \in \Xi$ and $x \in [0, \bar{x}]$, $F_\xi(-\infty \mid \xi, (\alpha, x)) = 1 - \alpha$. In what follows, we will write $F_\xi(\xi' \mid \xi, x)$ instead of $F_\xi(\xi' \mid \xi, (1, x))$. We assume that the map $x \mapsto F_\xi(\cdot \mid \xi, x)$ is continuous for all $\xi \in \Xi$.[5]

[3] Rather than imposing this bound, one could strengthen Assumption 1 by requiring that the marginal cost of investment grow without bound, i.e., $\partial c(x,\nu)/\partial x \to \infty$ as $x \to \infty$.

[4] We restrict attention to scalar firm characteristics ξ_{it} and investment x_{it} for clarity of exposition. Extending the analysis to multidimensional characteristics and actions would clutter notation and raise computational complexity. Few papers in the literature estimate or study dynamic games with multiple continuous choices; for a recent contribution in that direction, see Wang and Zhai (2025).

[5] As we will soon restrict ourselves to finite Ξ, we do not specify the topology in the space of conditional distributions $F_\xi(\cdot \mid \xi, x)$.

Assumption CINS rules out spillover effects in the evolution of firms' characteristics.[6] This excludes some models of interest from the subsequent analysis. For instance, models of dynamic price competition with persistence in demand due to switching costs such as Egan et al. (2025) do not satisfy this assumption. The endogenous state variable in these models is the vector of firms' lagged market shares. Firms' lagged market shares are not conditionally independent and are subject to spillovers, as market shares depend on all prices. Many other models of interest, however, do satisfy Assumption CINS, such as models of capacity accumulation (Besanko and Doraszelski (2004), Ryan (2012)) and quality ladder models (Doraszelski and Pakes (2007), Hashmi and van Biesebroeck (2016)).

2.1 Strategies

The assumptions we have made imply *ex-ante* firm symmetry, so we restrict attention to Symmetric Markov Perfect Equilibria (SMPE), as is standard in the literature. The Markov restriction constrains firm behavior to depend only on payoff-relevant variables: publicly observed firm characteristics $\boldsymbol{\xi}_t$ and private information $\varepsilon_{it} = (\rho_{it}, \kappa_{it}, \nu_{it})$. Symmetry imposes on value and policy functions the analogues of (1) and (2): under (1), it suffices to solve firm 1's dynamic programming problem; under (2), we can further reduce the state space from $\boldsymbol{\Xi} := \Xi^{\bar{N}}$ to $\boldsymbol{\Xi}^R := \{\boldsymbol{\xi} \in \boldsymbol{\Xi} : \xi_2 \leq \xi_3 \leq \ldots \leq \xi_{\bar{N}}\}$, by identifying states that are equivalent from firm 1's perspective. Given a state $\boldsymbol{\xi} \in \boldsymbol{\Xi}^R$, we denote by $\boldsymbol{\xi}^{(j)}$ the corresponding state from firm j's perspective, i.e., $\boldsymbol{\xi}^{(j)} = (\xi_j, s(\boldsymbol{\xi}_{-j}))$, where $s(\boldsymbol{\xi}_{-j})$ is $\boldsymbol{\xi}_{-j}$ sorted in increasing order.

In what follows, we will denote a strategy by

$$\sigma(\boldsymbol{\xi}, \varepsilon) = \left(\alpha^I(\boldsymbol{\xi}, \rho), \alpha^E(\boldsymbol{\xi}, \kappa), \sigma^x(\boldsymbol{\xi}, \nu)\right) ,$$

where $\varepsilon = (\rho, \kappa, \nu)$ and α^I and α^E denote, respectively, the incumbent's and entrant's decision to be active in $t + 1$.

[6] Assumption CINS also abstracts from aggregate shocks that affect transitions of all firms simultaneously. Such shocks — capturing, for instance, exogenous improvements to the outside option — can be accommodated at the cost of additional notation, and our uniqueness and existence results (Proposition 1 and Proposition 2, respectively) continue to hold in that setting.

2.2 The Incumbent's Problem

In what follows, we denote by $V_I(\boldsymbol{\xi}, \rho)$ the expected net present value (ENPV) of an incumbent faced with public state $\boldsymbol{\xi}$ and scrap value ρ and by $V_I^A(\boldsymbol{\xi}, \nu)$ the ENPV of an incumbent that has chosen to be active and has observed investment cost shock ν.

2.2.1 An Active Incumbent's Investment Problem

Fix a strategy σ. The objects $V_I^A(\boldsymbol{\xi}, \nu)$ and $V_I(\boldsymbol{\xi}, \rho)$ are related through the Bellman equation

$$V_I^A(\boldsymbol{\xi}, \nu) = \max_{x \in [0, \bar{x}]} \left\{ \pi(\boldsymbol{\xi}) - c(x, \nu) + \beta \mathbb{E}^\sigma [V_I(\boldsymbol{\xi}', \rho) \mid \boldsymbol{\xi}, x] \right\} \tag{4}$$

where

$$\mathbb{E}^\sigma [V_I(\boldsymbol{\xi}', \rho) \mid \boldsymbol{\xi}, x] = \int_{\boldsymbol{\varepsilon}_{-1}} \int_{\boldsymbol{\xi}'} \int_{\rho} V_I(\boldsymbol{\xi}', \rho) \, \mathrm{d}F_\rho \, \mathrm{d}F(\boldsymbol{\xi}' \mid \boldsymbol{\xi}, (1, x), \sigma_{-1}(\boldsymbol{\xi}, \boldsymbol{\varepsilon}_{-1})) \, \mathrm{d}G_{\boldsymbol{\varepsilon}_{-1}} \tag{5}$$

and $\sigma_{-1}(\boldsymbol{\xi}, \boldsymbol{\varepsilon}_{-1}) = (\sigma(\boldsymbol{\xi}^{(2)}, \varepsilon_2), \dots, \sigma(\boldsymbol{\xi}^{(\bar{N})}, \varepsilon_{\bar{N}}))$.

Let $\bar{V}_I(\boldsymbol{\xi}) \coloneqq \int_\rho V_I(\boldsymbol{\xi}, \rho) \, \mathrm{d}F_\rho$. Under Assumption CINS and the independence of the private shocks ε_i across firms, the expectation in (5) can be written as

$$\mathbb{E}^\sigma [V_I(\boldsymbol{\xi}', \rho) \mid \boldsymbol{\xi}, x] = \int_{\xi_1'} W(\xi_1' \mid \boldsymbol{\xi}; F^\sigma) \, \mathrm{d}F(\xi_1' \mid \xi_1, x) \, , \tag{6}$$

where

$$F^\sigma(\xi_j' \mid \boldsymbol{\xi}^{(j)}) \coloneqq \int_{\varepsilon_j} F(\xi_j' \mid \xi_j, \sigma(\boldsymbol{\xi}^{(j)}, \varepsilon_j)) \, \mathrm{d}G_{\varepsilon_j} \tag{7}$$

is the σ-induced marginal transition for firm j, and

$$W(\xi_1' \mid \boldsymbol{\xi}; F^\sigma) \coloneqq \int_{\xi_2'} \dots \int_{\xi_{\bar{N}}'} \bar{V}_I(\xi_1', \boldsymbol{\xi}_{-1}') \, \mathrm{d}F^\sigma(\xi_{\bar{N}}' \mid \boldsymbol{\xi}^{(\bar{N})}) \dots \mathrm{d}F^\sigma(\xi_2' \mid \boldsymbol{\xi}^{(2)}) \tag{8}$$

is the ENPV of an incumbent that reaches characteristic ξ_1' when competitors' characteristics evolve according to F^σ. Appendix A.1 derives equation (6).

Uniqueness of the Investment Decision. The first-order condition of the maximization problem on the right-hand side of (4) is

$$-\frac{\partial c(x,\nu)}{\partial x} + \beta \frac{\partial}{\partial x}\left(\int_{\xi_1'} W(\xi_1' \mid \boldsymbol{\xi}; F^\sigma)\, \mathrm{d}F(\xi_1' \mid \xi_1, x)\right) \leq 0\ ,$$

with equality if the solution is interior. It is desirable for the investment first-order condition to be sufficient for an optimum. Sufficiency is useful both in equilibrium computation and in estimation based on firms' optimal policies. One approach to establish sufficiency of the investment first-order condition is to adapt the concept of UIC-admissible transitions of Doraszelski and Satterthwaite (2010) to the transition process described above.

Definition 1 (Doraszelski and Satterthwaite (2010))**.** The distribution $F(\xi' \mid \xi, x)$ is UIC-admissible if, for all $\xi', \xi \in \Xi$ and $x \in [0, \bar{x}]$,

$$F(\xi' \mid \xi, x) = L(\xi', \xi) + K(\xi', \xi)Q(\xi, x)\ , \tag{9}$$

where $Q(\xi, x)$ is twice-continuously differentiable in x for all $\xi \in \Xi$ and $x \in [0, \bar{x}]$, and satisfies

$$\frac{\partial Q(\xi,x)}{\partial x} > 0 \quad \text{and} \quad \frac{\partial^2 Q(\xi,x)}{\partial x^2} < 0\ .$$

Proposition 1. *If Assumption 1 holds and $F(\xi' \mid \xi, x)$ is UIC-admissible, then the problem*

$$\max_{x\in[0,\bar{x}]} \pi(\boldsymbol{\xi}) - c(x,\nu) + \beta \int_{\xi_i'} W(\xi_1' \mid \boldsymbol{\xi}; F^\sigma)\, \mathrm{d}F(\xi_1' \mid \xi_1, x)$$

has a unique solution for all $\boldsymbol{\xi}$ and ν. Moreover, the maximizer is increasing in ν, and strictly so over the range of ν where it is interior.

Proof. See Appendix A.2. The proof is a rewriting of the proof of Proposition 3 in Doraszelski and Satterthwaite (2010) specialized to the transition process described above. □

Some transition processes of interest do not satisfy UIC-admissibility. For instance, in capacity accumulation games with deterministic transitions and no depreciation, the transition function is given by $F(\xi' \mid \xi, x) = \mathbf{1}\{\xi' \geq \xi + x\}$, which is not UIC-admissible. The decomposition in condition (9) says that transitions can be decomposed into a baseline distribution over ξ', the L term, a factor that spreads mass over ξ', the K term, and a factor that determines the mass to be spread over ξ', the Q term. Thus, the way in which a firm's current efforts to improve its state affect its state tomorrow is set by the K factor, and outside the firm's direct control. Models in which the firm directly controls the distribution of its future state, such as the capacity accumulation game described above, violate UIC-admissibility.

We provide an alternative condition for the uniqueness and monotonicity of the investment decisions in Proposition 3 in Appendix A.2. This condition places alternative restrictions on the transtition process and avoids the decomposition (9), at the cost of assuming that in equilibrium the W function is increasing in the firm's state ξ_1'. In the context of quality ladder or capacity accumulation games, this says that higher quality or capacity levels are beneficial for the firm. In applications where this assumption is deemed reasonable, Proposition 3 expands the range of applicability of the methods described in this paper. In particular, capacity accumulation games where firms' future capacity levels are stochastic but firms control the location of that distribution violate UIC-admissibility but satisfy the conditions of Proposition 3.[7]

2.2.2 The Incumbent's Exit Decision

The incumbent commits to an exit decision before observing the investment cost shock ν. Therefore, it must make its decision on the basis of its expected continuation value conditional on being active:

$$\bar{V}_I^A(\boldsymbol{\xi}) := \int V_I^A(\boldsymbol{\xi}, \nu)\, \mathrm{d}F_\nu \ . \tag{10}$$

Recall $V_I(\boldsymbol{\xi}, \rho)$ denotes the ENPV of an incumbent with scrap value ρ. Then

$$V_I(\boldsymbol{\xi}, \rho) = \max\left\{\pi(\boldsymbol{\xi}) + \rho, \bar{V}_I^A(\boldsymbol{\xi})\right\} = \max_{\chi\in\{0,1\}} \chi \bar{V}_I^A(\boldsymbol{\xi}) + (1-\chi)[\pi(\boldsymbol{\xi}) + \rho] \ . \tag{11}$$

The implied conditional probability of an incumbent remaining active is

$$\mathbb{P}\big(\alpha^I(\boldsymbol{\xi}, \rho) = 1 \mid \boldsymbol{\xi}\big) = F_\rho\big(\bar{V}_I^A(\boldsymbol{\xi}) - \pi(\boldsymbol{\xi})\big) \ . \tag{12}$$

[7]Proposition 3 requires transitions to be C^2 in the firm's control variable, and thus does not cover deterministic capacity accumulation games. Uniqueness of the investment choice in those models seems to require assumptions about the curvature of W.

2.3 The Entrant's Problem

Denote by $V_E^A(\boldsymbol{\xi}_{-1}, \nu)$ the ENPV of a potential entrant that enters under public state $\boldsymbol{\xi}_{-1}$ and draws investment cost shock ν. This function is characterized by

$$V_E^A(\boldsymbol{\xi}_{-1}, \nu) = \max_{x \in [0,\bar{x}]} \left\{ -c(x, \nu) + \beta \int_{\xi_1'} W(\xi_1' \mid (-\infty, \boldsymbol{\xi}_{-1}); F^\sigma) \, \mathrm{d}F(\xi_1' \mid \xi_e, x) \right\} , \tag{13}$$

where $\xi_e \in \Xi$ is an exogenously specified initial quality level for potential entrants.

Potential entrants either enter the market or perish. Therefore, their ENPV given entry cost κ is

$$V_E(\boldsymbol{\xi}_{-1}, \kappa) = \max \left\{ 0, \bar{V}_E^A(\boldsymbol{\xi}_{-1}) - \kappa \right\} = \max_{\chi \in \{0,1\}} \chi \big[\bar{V}_E^A(\boldsymbol{\xi}_{-1}) - \kappa \big] \tag{14}$$

where we have normalized the value of entrants' outside option to zero and

$$\bar{V}_E^A(\boldsymbol{\xi}_{-1}) := \int V_E^A(\boldsymbol{\xi}_{-1}, \nu) \, \mathrm{d}F_\nu \, . \tag{15}$$

The conditional probability of entry is

$$\mathbb{P}\big(\alpha^E(\boldsymbol{\xi}_{-1}, \kappa) = 1 \mid \boldsymbol{\xi}_{-1}\big) = F_\kappa\big(\bar{V}_E^A(\boldsymbol{\xi}_{-1})\big) \, . \tag{16}$$

2.4 Equilibrium

Definition 2. Let $\Xi_I := (\Xi \setminus \{-\infty\}) \times \Xi^{\bar{N}-1}$. A Symmetric Markov Perfect Equilibrium (SMPE) is a pair $(\bar{V}_I, \sigma)$ where $\bar{V}_I : \Xi_I \to \mathbb{R}$ and $\sigma = (\sigma^x, \alpha^E, \alpha^I)$ are such that

1. $\sigma^x(\boldsymbol{\xi}, \nu)$ solves problem (4) subject to (6), (7), and (8), for all $\boldsymbol{\xi} \in \Xi_I$ and ν in the support of F_ν.

2. $\alpha^I(\boldsymbol{\xi}, \rho)$ solves problem (11) subject to (10) and (4), for all $\boldsymbol{\xi} \in \Xi_I$ and ρ in the support of F_ρ.

3. $\alpha^E(\boldsymbol{\xi}_{-1}, \kappa)$ solves problem (14) subject to (15) and (13), for all $\boldsymbol{\xi}_{-1} \in \Xi^{\bar{N}-1}$ and κ in the support of F_κ.

4. For all $\boldsymbol{\xi} \in \Xi_I$,

$$\begin{aligned}\bar{V}_I(\boldsymbol{\xi}) &= \int \left\{\alpha^I(\boldsymbol{\xi},\rho)\bar{V}_I^A(\boldsymbol{\xi}) + \left(1-\alpha^I(\boldsymbol{\xi},\rho)\right)[\pi(\boldsymbol{\xi})+\rho]\right\} \mathrm{d}F_\rho \\ &= F_\rho\left(\bar{V}_I^A(\boldsymbol{\xi}) - \pi(\boldsymbol{\xi})\right)\bar{V}_I^A(\boldsymbol{\xi}) \\ &+ \left[1 - F_\rho\left(\bar{V}_I^A(\boldsymbol{\xi}) - \pi(\boldsymbol{\xi})\right)\right]\left(\pi(\boldsymbol{\xi}) + \mathbb{E}[\rho \mid \rho \geq \bar{V}_I^A(\boldsymbol{\xi}) - \pi(\boldsymbol{\xi})]\right), \end{aligned} \tag{17}$$

where $\bar{V}_I^A(\boldsymbol{\xi})$ is given by (10).

We define an equilibrium only in terms of incumbents' integrated value functions. This is due to the assumption, common in this literature, that firms that choose to be inactive in the following period perish. As a result, when a firm's quality becomes $\xi' = -\infty$ (as a consequence of exit or no entry), that firm's continuation value is zero. Therefore, $\bar{V}_I$ alone is sufficient to determine firm behavior.[8]

To the best of our knowledge, the existing proofs of existence of a Markov Perfect Equilibrium in Ericson and Pakes (1995) type models — notably, Doraszelski and Satterthwaite (2010) — do not allow for shocks to firms' costs of investment. Therefore, those proofs do not apply to the environment described above. We thus provide an existence result.

Proposition 2. Suppose $|\Xi| < \infty$ and transitions $F(\xi' \mid \xi, x)$ satisfy Definition 1. Then, a Symmetric Markov Perfect Equilibrium exists.

Proof. See Appendix A.3. □

The added complication brought about by investment cost shocks is that the investment policy is infinite-dimensional. Random scrap values and setup costs without investment cost shocks — as in Doraszelski and Satterthwaite (2010) — do not raise this difficulty. In those environments, entry and exit policies have a finite-dimensional representation as cutoffs or probabilities that depend only on the public state $\boldsymbol{\xi}$.

To establish Proposition 2, one potential strategy is to invoke a fixed-point theorem applicable to infinite-dimensional spaces. However, under the finiteness assumption of Proposition 2, the more elementary Brouwer fixed-point theorem suffices. The idea is to establish the existence of a fixed point of a certain map from the set of collections of probability distributions of the form

[8]Note also that we do not define $\bar{V}_I$ on $\boldsymbol{\Xi}^R$ but rather on $\boldsymbol{\Xi}_I$. Similarly we do not define the policy functions on the reduced state space. This is for the sake of precision, as strategies must be complete contingent plans. However, the discussion in Section 2.1 applies. In particular, when computing equilibria we do exploit symmetry, as we discuss shortly.

$\{F^\sigma(\xi' \mid \boldsymbol{\xi}) : \xi' \in \Xi, \boldsymbol{\xi} \in \boldsymbol{\Xi}^R\}$ into itself. These collections of probability distributions are finite-dimensional objects under the finiteness assumption of Proposition 2, which enables the application of Brouwer's theorem. A symmetric Markov Perfect Equilibrium is then constructed using the strategy profile in which each firm plays the unique optimal policy subject to competitors' qualities having the transition described by the fixed-point. We provide further details in Appendix A.3.

This proof strategy informs our computation of Symmetric Markov Perfect Equilibria, in that in our iterative procedure we check for convergence of the parameters of firms' optimization problems, namely $\bar{V}_I$ and F^σ, rather than of the policy functions themselves. We start with guesses for $\bar{V}_I(\boldsymbol{\xi})$ and $F^\sigma(\xi' \mid \boldsymbol{\xi})$. With these two objects we can compute $W(\xi' \mid \boldsymbol{\xi}; F^\sigma)$. We then solve for firms' optimal investment and entry and exit decisions, i.e., we perform the computations associated with conditions 1 to 3 in Definition 2. These computations yield $\bar{V}_I^A(\boldsymbol{\xi})$ and allow us to update F^σ.[9] We then update $\bar{V}_I(\boldsymbol{\xi})$ using condition 4 in Definition 2.[10] We iterate on these steps until both $\bar{V}_I$ and F^σ converge. As the firms' objective function depends solely and continuously on these objects, Berge (1963)'s Maximum Theorem implies that the policy functions associated with the firms' problem also converge. By the symmetry assumption, it is sufficient to compute $\bar{V}_I(\boldsymbol{\xi})$ and $F^\sigma(\xi' \mid \boldsymbol{\xi})$ for $\boldsymbol{\xi} \in \boldsymbol{\Xi}^R$. As shown in Pakes and McGuire (1994), the reduced state space $\boldsymbol{\Xi}^R$ grows in the number of firms as a polynomial of order $|\Xi|$ rather than exponentially. For instance, symmetry reduces the state space cardinality of the Hashmi and van Biesebroeck (2016) model we consider from 1,048,576 to 62,016.

3 Estimators

This section introduces a pseudo-MLE estimator based on recursive equilibrium conditions in subsection 3.1, where we also discuss alternative estimators for models where the pseudo-MLE does not apply. Subsection 3.2 discusses the estimation of firms' integrated value functions, a key input to all estimators of dynamic games.

[9] Updating F^σ involves an integral with respect to F_ν — see equation (40) in Appendix A.3. The choice of approximant to that integral will determine the values of ν for which we compute firms' optimal investment choices.

[10] For suitable distributions, e.g. the uniform, exponential, and lognormal, the expectation in that condition can be written in closed form.

3.1 A Pseudo MLE Estimator Based on Recursive Equilibrium Conditions

This subsection presents estimators based on optimality conditions 1–3 in Definition 2. Let the investment cost function depend on parameters $\boldsymbol{\theta}_x$, the distribution of scrap values depend on parameters $\boldsymbol{\theta}_\rho$, and the distribution of entry costs depend on parameters $\boldsymbol{\theta}_\kappa$. We denote $\boldsymbol{\theta}_{-\kappa} := (\boldsymbol{\theta}_x, \boldsymbol{\theta}_\rho)$ and, collectively, $\boldsymbol{\theta} := (\boldsymbol{\theta}_{-\kappa}, \boldsymbol{\theta}_\kappa)$. The econometrician aims to estimate $\boldsymbol{\theta}$.

Suppose we can obtain an estimate (up to parameters) $\widehat{\bar{V}}_I(\boldsymbol{\xi}; \boldsymbol{\theta}_{-\kappa})$ of the ex-ante value function $\bar{V}_I(\boldsymbol{\xi})$.[11] Suppose we also have estimates $\hat{F}(\xi' \mid \xi, x)$ of the quality transitions $F(\xi' \mid \xi, x)$ and $\hat{F}^\sigma(\xi' \mid \boldsymbol{\xi})$ of the integrated equilibrium transition probabilities defined in equation (7). These estimates allow us to set up an empirical analog of firms' investment problem:

$$\max_{x \in [0, \bar{x}]} \left\{ \pi(\boldsymbol{\xi}) - c(x, \nu; \boldsymbol{\theta}_x) + \beta \int_{\xi_1'} \widehat{W}(\xi_1' \mid \boldsymbol{\xi}; \hat{F}^\sigma, \boldsymbol{\theta}_{-\kappa}) \, \mathrm{d}\hat{F}(\xi_1' \mid \xi_1, x) \right\}, \tag{18}$$

where $\widehat{W}(\xi_1' \mid \boldsymbol{\xi}; \hat{F}^\sigma, \boldsymbol{\theta}_{-\kappa})$ is given by (8) substituting $\widehat{\bar{V}}_I(\boldsymbol{\xi}; \boldsymbol{\theta}_{-\kappa})$ for $\bar{V}_I(\boldsymbol{\xi})$ and $\hat{F}^\sigma$ for F^σ. We base estimation on necessary conditions for observed investment levels to solve (18) and on empirical analogs to equations (12) and (16). We sketch the main ideas below, and save full details for Appendix B.

Let x be an observed investment level. The conditional distribution of investment given $\boldsymbol{\xi}$ is given by

$$\begin{aligned} F_X(x \mid \boldsymbol{\xi}) &= \mathbb{P}(\sigma^x(\boldsymbol{\xi}, \nu) \leq x \mid \boldsymbol{\xi}) \\ &= \mathbb{P}\left(\nu \geq (\sigma^x)^{-1}(x; \boldsymbol{\xi}) \mid \boldsymbol{\xi}\right) \\ &= 1 - F_\nu\left((\sigma^x)^{-1}(x; \boldsymbol{\xi})\right), \end{aligned} \tag{19}$$

where $(\sigma^x)^{-1}(x; \boldsymbol{\xi})$ is the inverse of the investment policy with respect to ν.[12] The second equality uses the monotonicity of the investment policy in ν, sufficient conditions for which are given in Proposition 1 and Proposition 3. When $x = 0$, equation (19) gives the investment contribution to the likelihood. When $x > 0$, we differentiate (19) to obtain the conditional density $f_X(x \mid \boldsymbol{\xi})$. We

[11] We discuss alternative estimators of $\bar{V}_I(\boldsymbol{\xi})$ in Section 3.2. Observe that our notation indicates that these value-function estimates do not depend on $\boldsymbol{\theta}_\kappa$. We discuss why below.

[12] When $x = 0$ we define $(\sigma^x)^{-1}(0; \boldsymbol{\xi}) := \inf\{\nu \in \operatorname{supp}(F_\nu) : \sigma^x(\boldsymbol{\xi}, \nu) = 0\}$.

obtain

$$f_X(x \mid \boldsymbol{\xi}) = -f_\nu \left((\sigma^x)^{-1}(x;\boldsymbol{\xi})\right) \cdot \frac{\partial}{\partial x}(\sigma^x)^{-1}(x;\boldsymbol{\xi}) \tag{20}$$

$$= -f_\nu \left((\sigma^x)^{-1}(x;\boldsymbol{\xi})\right) \times \frac{-\partial_x^2 c(x,(\sigma^x)^{-1}(x;\boldsymbol{\xi})) + \partial_x MB(\boldsymbol{\xi},x)}{\partial_{\nu x}^2 c(x,(\sigma^x)^{-1}(x;\boldsymbol{\xi}))} , \tag{21}$$

where $MB(\boldsymbol{\xi},x) := \partial_x \left(\beta \int_{\xi_1'} \widehat{W}(\xi_1' \mid \boldsymbol{\xi}; \hat{F}^\sigma, \boldsymbol{\theta}_{-\kappa})\, \mathrm{d}\hat{F}(\xi_1' \mid \xi_1, x)\right)$. Appendix B gives additional details for the case when $F(\xi' \mid \xi, x)$ is UIC-admissible.

We now turn to exit decisions. If a is an indicator for whether an incumbent chooses to be active, then the contribution to the likelihood is given by

$$aF_\rho \left(\widehat{\hat{V}}_I^A(\boldsymbol{\xi};\boldsymbol{\theta}_{-\kappa}) - \pi(\boldsymbol{\xi}); \boldsymbol{\theta}_\rho\right) + (1-a)\left[1 - F_\rho \left(\widehat{\hat{V}}_I^A(\boldsymbol{\xi};\boldsymbol{\theta}_{-\kappa}) - \pi(\boldsymbol{\xi}); \boldsymbol{\theta}_\rho\right)\right] ,$$

where $\widehat{\hat{V}}_I^A(\boldsymbol{\xi};\boldsymbol{\theta}_{-\kappa})$ is an estimate of the value of being active.

There are two approaches to computing the estimate of the value of being active. They trade off computational cost and efficiency. The first approach makes use of an estimate of the policy function, $\hat{\sigma}^x(\boldsymbol{\xi},\nu)$. With such an estimate in hand, we can estimate an incumbent's value of being active as

$$\widehat{\hat{V}}_I^A(\boldsymbol{\xi};\boldsymbol{\theta}_{-\kappa}) = \int \hat{V}_I^A(\boldsymbol{\xi},\nu;\hat{\sigma}^x,\boldsymbol{\theta}_{-\kappa})\, \mathrm{d}F_\nu ,$$

where $\hat{V}_I^A(\boldsymbol{\xi},\nu;\hat{\sigma}^x,\boldsymbol{\theta}_{-\kappa})$ substitutes $\hat{\sigma}^x(\boldsymbol{\xi},\nu)$ for x in the objective function of problem (18). This approach avoids solving that maximization problem and is thus computationally cheap. The second approach consists in estimating $V_I^A(\boldsymbol{\xi},\nu)$ by solving the maximization problem in (18) at pre-specified values of ν to approximate the integral with respect to F_ν. This approach is computationally more demanding, but uses information regarding how parameters affect optimal investment decisions, and is thus more efficient. In our Monte Carlo experiments in Section 4, we implement the first approach and report satisfactory results. The likelihood contribution of entry decisions is analogous to the contribution of exit decisions and omitted for brevity.

Let $L_X(\boldsymbol{\theta};x,\boldsymbol{\xi})$ and $L_A(\boldsymbol{\theta};a,\boldsymbol{\xi})$ denote the contributions to the likelihood stemming, respectively, from firms' investment choices and decisions to be active. The likelihood is then

$$L(\boldsymbol{\theta}) = \prod_{j,m} L_X(\boldsymbol{\theta};x_{jm},\boldsymbol{\xi}_{jm}) \times L_A(\boldsymbol{\theta};a_{jm},\boldsymbol{\xi}_{jm}) . \tag{22}$$

Evaluating expression (22) is computationally inexpensive. The most expensive step is computing $\widehat{W}(\xi_1' \mid \boldsymbol{\xi}; \hat{F}^\sigma, \boldsymbol{\theta}_{-\kappa})$, an average of the value function estimates $\widehat{\hat{V}}_I(\boldsymbol{\xi}; \boldsymbol{\theta}_{-\kappa})$. As shown in Section 3.2, computing the value function estimates reduces to a matrix-vector multiplication. The inverse policy $(\sigma^x)^{-1}$ and the cost derivatives $\partial_x^2 c$ and $\partial_{\nu x} c$ are available in closed form, given the assumed functional form of the investment cost function. Computing the derivative of the marginal benefit function, $\partial_x MB(\boldsymbol{\xi}, x)$, amounts to a one-dimensional integral approximation and, in the leading UIC-admissible case, simplifies further to calculating a simpler integral of $\widehat{W}$ terms and the derivative of the transition function's Q factor, which has an assumed functional form (see Appendix B). The exit and entry contributions are likewise cheap to compute. They require only the values of being active, which under the first approach discussed above are obtained by substituting the estimated policy into (18), avoiding any optimization.

An alternative estimator There are alternative ways to use firms' optimality conditions for estimation. For instance, in Section 4.1 we use a nonlinear least squares estimator to estimate the dynamic game in Bajari et al. (2007). This estimator is based on minimizing squared deviations between observed investment levels and investment levels predicted by the model, as well as squared deviations between observed decisions to be active and their predicted probabilities. Model-predicted investment levels are obtained by solving the investment problem in (18). As above, predicted probabilities of being active can be obtained either by using an estimate of the investment policy, or by using the model-predicted optimal level of investment. We use the former approach.

The results in Section 4.1 show that this estimator works well. We use it to estimate the model in Bajari et al. (2007) because that model does not have shocks to the cost of investment, making the pseudo maximum likelihood (PMLE) estimator above inapplicable. For models with investment cost shocks, we favor the PMLE estimator, as it is significantly cheaper to compute.[13]

3.2 Estimating Integrated Value Functions

Integrated value functions are a key input to the pseudo-MLE and NLLS estimators discussed above, as they are to the Bajari et al. (2007) estimator. This

[13] One could use the solutions to problem (18) to form other estimators. For instance, we have experimented with an indirect inference estimator and found that it also works well.

subsection discusses methods to estimate them up to parameters.

Bajari et al. (2007) propose estimating integrated value functions by forward simulation. As they note, linearity of the value function with respect to $\boldsymbol{\theta}$, which is a feature of many models including those in Section 4, significantly reduces the computational burden of forward simulation. If $\bar{V}(\boldsymbol{\xi}; \sigma', \sigma, \boldsymbol{\theta})$ denotes the value for a firm of being in state $\boldsymbol{\xi}$ when its policy is σ' and its competitors all play the policy σ, then there exists a function $\bar{\Lambda}(\boldsymbol{\xi}, \sigma', \sigma)$ such that $\bar{V}(\boldsymbol{\xi}; \sigma', \sigma, \boldsymbol{\theta}) = \bar{\Lambda}(\boldsymbol{\xi}, \sigma', \sigma) \cdot \boldsymbol{\theta}$ and forward simulation need not be repeated as $\boldsymbol{\theta}$ varies.

Alternatively, as we now show, one can solve for $\bar{\boldsymbol{V}}_I(\boldsymbol{\xi}; \boldsymbol{\theta}_{-\kappa})$ in closed form when the state space is finite. Let $P(\xi' \mid \xi, a)$ denote the probability that a firm's quality in $t+1$ is ξ' conditional on its current quality being ξ and its action being $a = (\alpha, x)$. Moreover, let Ξ_I^R denote the set of states in the reduced state space in which firm 1 is active, i.e., $\Xi_I^R := \{\boldsymbol{\xi} \in \Xi^R : \xi_1 > -\infty\}$. Let $\bar{\boldsymbol{V}}_I = [\bar{V}_I(\boldsymbol{\xi}) : \boldsymbol{\xi} \in \Xi_I^R]$ be a vector stacking incumbents' integrated values across states in Ξ_I^R. We show in Appendix A.4 that $\bar{\boldsymbol{V}}_I$ satisfies

$$[I - \beta \boldsymbol{M}(\boldsymbol{P})]\bar{\boldsymbol{V}}_I = \boldsymbol{\pi} - \boldsymbol{K}(\boldsymbol{\theta}_x) + \boldsymbol{\Sigma}(F_\rho) \tag{23}$$

where

$$\boldsymbol{K}(\boldsymbol{\theta}_x) = \left[\mathbb{P}_I^A(\boldsymbol{\xi}) \int c(\sigma^x(\boldsymbol{\xi}, \nu), \nu; \boldsymbol{\theta}_x)\, \mathrm{d}F_\nu \;:\; \boldsymbol{\xi} \in \Xi_I^R\right] \tag{24}$$

$$\boldsymbol{\Sigma}(F_\rho) = \left[[1 - \mathbb{P}_I^A(\boldsymbol{\xi})]\mathbb{E}[\rho \mid \rho > F_\rho^{-1}(\mathbb{P}_I^A(\boldsymbol{\xi})); \boldsymbol{\theta}_\rho] \;:\; \boldsymbol{\xi} \in \Xi_I^R\right] \tag{25}$$

where $\mathbb{P}_I^A(\boldsymbol{\xi}) = \mathbb{P}(\alpha^I(\boldsymbol{\xi}, \rho) = 1 \mid \boldsymbol{\xi})$ is the probability that the incumbent chooses to be active in state $\boldsymbol{\xi}$, and $\boldsymbol{M}(\boldsymbol{P})$ is the transition matrix implied by the policy function σ, i.e.,[14]

$$\boldsymbol{M}(\boldsymbol{P}) = \left[\mathbb{P}^\sigma(\boldsymbol{\xi}_l \mid \boldsymbol{\xi}_k) \;:\; 1 \leq l, k \leq |\Xi_I^R|\right] \tag{26}$$

where

$$\mathbb{P}^\sigma(\boldsymbol{\xi}' \mid \boldsymbol{\xi}) = \prod_{j=1}^{\bar{N}} P^\sigma(\xi_j' \mid \boldsymbol{\xi}) = \prod_{j=1}^{\bar{N}} \int P(\xi_j' \mid \xi_j, \sigma(\boldsymbol{\xi}^{(j)}, \varepsilon))\, \mathrm{d}G_\varepsilon\,. \tag{27}$$

Equations (23)–(27) imply that we can estimate $\bar{\boldsymbol{V}}_I$ up to parameters by es-

[14]We order states in Ξ^R and Ξ_I^R lexicographically, where firm 1 takes precedence over firm 2, which takes precedence over firm 3, and so on. We do so by interpreting $\boldsymbol{\xi} = (\xi_1, \ldots, \xi_{\bar{N}})$ as a number in base $|\Xi|$.

timating the probabilities $\mathbb{P}_I^A(\boldsymbol{\xi})$ and $P^\sigma(\xi' \mid \boldsymbol{\xi})$ and the investment policy function $\sigma^x(\boldsymbol{\xi}, \nu)$. Since firms' characteristics are observable, these probabilities are directly estimable from the data. For instance, if the data allow, one could estimate these probabilities using frequency estimators. We discuss estimation of the investment policy function below.

Equation (23) allows us to solve efficiently for $\bar{\boldsymbol{V}}_I(\boldsymbol{\xi}; \boldsymbol{\theta}_{-\kappa})$. That system of equations can essentially be solved only once, as we can compute a decomposition of $I - \beta \boldsymbol{M}(\boldsymbol{P})$ and store it in memory. Then, as we vary $\boldsymbol{\theta}_{-\kappa}$, we only need to recompute the expected flow profits and solve the resulting system using the stored matrix decomposition.

Much like with forward simulation, linearity of flow payoffs is computationally helpful, as can be readily seen from equation (23). If the cost of investment is linear in parameters and the scrap value is either non-existent (as in Hashmi & van Biesebroeck, 2016), deterministic (as in Bajari et al., 2007), or exponentially distributed, equation (23) implies that $\bar{V}_I$ is linear in parameters, i.e., $\bar{V}_I = [I - \beta \boldsymbol{M}(\boldsymbol{P})]^{-1} \boldsymbol{X} \boldsymbol{\theta}_{-\kappa}$ for some matrix $\boldsymbol{X}$. In this case there are additional computational savings, as one can store the solution A to $[I - \beta \boldsymbol{M}(\boldsymbol{P})]A = \boldsymbol{X}$ in memory and only compute $A\boldsymbol{\theta}_{-\kappa}$ as the parameters change.[15]

The trade-off between forward simulation and solving equation (23) is one between computational cost and simulation error. Forward simulation can be computationally cheaper when the state space is large, but involves simulation error that the closed-form solution avoids.[16]

Estimation of the investment policy function. To estimate $\sigma^x(\boldsymbol{\xi}, \nu)$, we build on Bajari et al. (2007). Their argument implies, in the case in which $\sigma^x(\boldsymbol{\xi}, \nu)$ is decreasing in ν, that

$$\sigma^x(\boldsymbol{\xi}, \nu) = F_X^{-1}(1 - F_\nu(\nu) \mid \boldsymbol{\xi}) \ , \tag{28}$$

where $F_X(x \mid \boldsymbol{\xi})$ is the distribution of investment conditional on $\boldsymbol{\xi}$, which is identified. That is, the policy function is identified by the quantiles of the condi-

[15]When these conditions on scrap values fail, equation (23) shows that $\bar{V}_I$ is no longer linear in parameters, even when the cost of investment is. The non-linearity arising from exit behavior does not, however, add substantive computational burden. In the closed-form approach, as parameters change one needs only to recompute the right-hand side of equation (23) and solve the linear system using the pre-computed matrix factorization. An analogous argument under forward simulation is given in Appendix D.2.

[16]Note that when using forward simulation, one can directly estimate $\hat{W}(\xi_1' \mid \boldsymbol{\xi}; \hat{F}^\sigma, \boldsymbol{\theta}_{-\kappa})$, bypassing the integrated value function.

tional distribution of investment.[17] Using this, the integral in equation (24) can be approximated by $\sum_{z=1}^{Z} \omega_z c(F_X^{-1}(1 - F_\nu(\nu_z) \mid \boldsymbol{\xi}), \nu_z; \boldsymbol{\theta}_x)$, for judiciously chosen weights ω_z and nodes ν_z. We estimate $F_X^{-1}(1 - F_\nu(\nu_i) \mid \boldsymbol{\xi})$ as the predicted values from quantile regressions of investment on features of $\boldsymbol{\xi}$. Appendix C describes the full set of first-stage estimators.

4 Monte Carlo Simulations

We use two models to compare the performance of the estimators based on recursive equilibrium conditions introduced in Section 3 with that of the BBL estimator. First, we consider a model similar to the one simulated and estimated in Bajari et al. (2007). Then, we extend the model by Hashmi and van Biesebroeck (2016) to allow for entry and exit. Both models are private cases of the model presented in Section 2 and feature UIC-admissible transitions. In both designs, potential entrants enter at the lowest active quality level, i.e., $\xi_e = \min(\Xi \setminus \{-\infty\})$.

4.1 A Bajari et al. (2007) Inspired Model

There are $\bar{N}$ single-product firms in the market, N of which are active. We index a firm and its product by j. The quality of product j, denoted ξ_j, is an element of the set $\Xi = \{-\infty, -\ln 20, -\ln 19, \ldots, \ln 19, \ln 20\}$. When consumer i purchases product j, she obtains utility

$$u_{ij} = \gamma_0 \xi_j + \alpha p_j + \varepsilon_{ij} \ , \tag{29}$$

where p_j is the price of product j and ε_{ij} are iid Type 1 Extreme Value random variables. At the beginning of each period, incumbent firms compete in prices à la Nash-Bertrand. The pricing game has a unique Nash equilibrium (Caplin & Nalebuff, 1991). Therefore, firms' flow profits $\pi_j(\boldsymbol{\xi})$ are well-defined.

Firms' current investment affects their product quality in the following period. For $\xi \in \Xi \setminus \{-\infty\}$, let ξ^+ and ξ^- denote, respectively, the next higher and next lower elements of $\Xi \setminus \{-\infty\}$, when they exist. For interior ξ_j, the

[17] This argument rests on the maintained assumption that F_ν is known. If F_ν is known up to parameters $\boldsymbol{\theta}_\nu$, equation (28) must account for that dependence. This in turn implies that $\boldsymbol{K}(\boldsymbol{\theta}_x)$ in equation (24) also depends on $\boldsymbol{\theta}_\nu$. It can then be seen from equation (23) that the integrated value function ceases to be linear in the structural parameters even when the investment cost function is linear in $\boldsymbol{\theta}_x$ and the scrap value is deterministic (as in Bajari et al., 2007).

firm's quality can move up to ξ_j^+, remain at ξ_j, or move down to ξ_j^-. Quality only transitions to $-\infty$ as a result of exit. Each firm's quality is affected by two shocks, one positive and one negative, which are independent from one another and across firms. The negative shock lowers quality with exogenous probability $\delta \in (0,1)$. The positive shock instead increases quality by one step with probability $u(x) = \frac{\psi x}{1+\psi x}$. This can be interpreted as the probability that investment yields a product improvement.

Given independence of the shocks, for interior ξ the quality transition probabilities satisfy[18]

$$P(\xi' \mid \xi, x) = \begin{cases} \delta[1-u(x)] & \text{if } \xi' = \xi^- \\ 1-\delta-u(x)(1-2\delta) & \text{if } \xi' = \xi \\ (1-\delta)u(x) & \text{if } \xi' = \xi^+ \\ 0 & \text{otherwise} \end{cases}. \tag{30}$$

Given this structure, firms choose investment to maximize the present-discounted stream of profits. They balance expected higher product quality with an immediate cost of investment $c(x) = \theta_x x$. There is no shock to the cost of investment. It is easy to see that the transitions in (30) are UIC-admissible.[19]

At the beginning of the period, incumbents observe a private scrap value $\rho \sim U[22,23]$ and potential entrants privately observe an entry cost $\kappa \sim U[22,30]$, and simultaneously make entry and exit decisions. Those decisions are implemented at the end of the period. We set the maximum number of firms to $\bar{N} = 3$.

The model in this section differs from the Bajari et al. (2007) model in five aspects. First, we use the utility specification in (29), whereas the Bajari et al. (2007) specification is logarithmic in net income. We set the values of γ_0 and

[18] The probability that $\xi' = \xi$ at maximum (minimum) quality is defined to be the complement of the probability of ξ decreasing (increasing).

[19] The local structure of the transitions (30) can be exploited for computational gains. It implies that the transition matrix $\boldsymbol{M}(\boldsymbol{P})$ in equation (26) is banded. A banded matrix is a sparse matrix whose non-zero entries are confined to a band around the main diagonal. Importantly, the LU decomposition of a banded matrix has banded components. This implies substantial computational savings in both the computation of the LU decomposition and the substitutions used to compute the solution to the linear system (23). It is important for this observation that equation (23) refers only to incumbent states. The matrix of transitions over all states includes non-zero entries in its first few columns (i.e., those that pertain to transitions to $\xi = -\infty$), making its lower bandwidth large, and reducing the computational savings. The full computational savings are thus a consequence of both the local nature of transitions and the assumption that exiting firms perish.

Table 1: BBL Model Parameters

Parameter	**Value**
Data Structure	
Number of Quality Levels	39
Minimum Quality Level	-2.99573
Maximum Quality Level	2.99573
Maximum Number of Players	3
Number of Markets	100
Number of Periods	40
Model Parameters	
Discount Factor (β)	0.925
Demand System (γ_0, α)	[0.1, -0.25]
Marginal Cost (θ_c)	[1.09861, 0.0]
Quality Transition (δ, ψ)	[0.7, 7.0]
Investment Cost (θ_x)	1.0
Investment Cost Shock Distribution (F_ν)	Dirac(0.0)
Scrap Value Distribution (F_ρ)	Uniform(22.0, 23.0)
Entry Costs Distribution (F_κ)	Uniform(22.0, 30.0)
Forward Simulation Settings	
Number of Simulated Paths	250
Simulation Horizon	150

Parameterization of the BBL model used in Monte Carlo simulations.

α to approximate their specification. Second, in the Bajari et al. (2007) model all firms experience the same negative shock to their qualities, interpreted as a stochastic improvement to the outside option, whereas we model negative shocks as firm-specific and independent. This also leads us to define the set of possible qualities differently from Bajari et al. (2007). There, both the outside option and firms' product qualities are integers between zero and 20. Rather than being explicit about the quality of the outside option, we normalize it to zero and define Ξ to be the set of all possible quality differences relative to the outside option in the BBL specification.

Third, we use a stochastic scrap value upon exit, whereas Bajari et al. (2007) assume a deterministic scrap value. We do so because we found it challenging to compute an equilibrium to the model with a deterministic scrap value to a satisfactory degree of accuracy. This is intuitive. When the scrap value is deterministic, small changes in the value function and endogenous transition probabilities lead firms to change their exit decisions, leading to failure in convergence. This causes the iterative solver to fail to converge. Fourth, we allow all potential entrants to enter, whereas Bajari et al. (2007) allow only one entrant per period. Fifth, our parameter values differ from those in Bajari et al. (2007). This is because under their parameters we computed equilibria in which potential entrants always enter, and incumbents never exit. That is perhaps due to our different utility specification. We choose parameter values that approximate the entry and exit probabilities reported in the supplementary materials to Bajari et al. (2007).

4.1.1 Estimators

The lack of a shock to the cost of investment makes the pseudo-MLE estimator inapplicable. Therefore, we estimate the BBL model using a nonlinear least squares estimator. Explicitly, we solve

$$\min_{\boldsymbol{\theta}} \sum_{j,m,t} \left\{ \left(x_{jmt} - \sigma^x(\boldsymbol{\xi}_{jmt}; \boldsymbol{\theta}, \hat{\Phi}) \right)^2 + \left(a_{jmt} - \mathbb{P}\left(\alpha(\boldsymbol{\xi}, \zeta) = 1 \mid \boldsymbol{\xi}_{jmt}, \boldsymbol{\theta}, \hat{\Phi} \right) \right)^2 \right\} . \tag{31}$$

In this expression, x_{jmt} is firm j's investment choice in market m and period t, $\sigma^x(\boldsymbol{\xi}_{jmt}; \boldsymbol{\theta}, \hat{\Phi})$ denotes the optimal investment when the firm's state is $\boldsymbol{\xi}_{jmt}$, given parameters $\boldsymbol{\theta}$ and first-stage estimates $\hat{\Phi}$, a_{jmt} is firm j's decision to be active in market m and period t, and $\mathbb{P}(\alpha(\boldsymbol{\xi}, \zeta) = 1 \mid \boldsymbol{\xi}_{jmt}; \boldsymbol{\theta}, \hat{\Phi})$ is the model-implied probability that a firm chooses to be active when it's state is $\boldsymbol{\xi}_{jmt}$, given parameters $\boldsymbol{\theta}$ and first-stage estimates $\hat{\Phi}$. The variable ζ should be interpreted as either the scrap value or the entry cost, depending on whether firm j is an incumbent or a potential entrant in market m and period t, as determined by $\boldsymbol{\xi}_{jmt}$. Appendix C describes the construction of $\hat{\Phi}$.

The BBL estimator is presented in Appendix D.1. Our implementation of the estimator minimizes (54) with respect to $\boldsymbol{\theta}$.[20] As evidenced by the criterion, the BBL estimator requires the econometrician take a stance on the class of deviations used in estimation. We implement two variants that have appeared in the literature, which we term *additive* and *multiplicative* deviations, as well as an *asymptotic* variant we devised. These classes of deviations perturb estimated policies additively, multiplicatively, or by drawing alternative coefficients for policy functions from the asymptotic distribution of estimated policy function parameters. We discuss implementation of each in detail in Appendix D.2. We set the number of criterion inequalities to the number of unique states in the data. As shown in Tables 1 and 3, we forward simulate 250 independent histories for 150 periods to approximate value functions.[21]

4.1.2 Monte Carlo Results

We simulate 500 independent datasets from the model described in Section 4.1 and estimate the parameters of the model using the NLLS estimator defined in equation (31) and the three implementations of the BBL estimator described in section 4.1.1. We report the results in Table 2.[22]

Table 2 reports the true values of each parameter along with the mean and standard deviation of parameter estimates across Monte Carlo simulations. The column labeled "Recursive" refers to the NLLS estimator. The labels "Additive", "Multiplicative", and "Asymptotic" refer to type of deviations used in each of the three implementations of the BBL estimator.

The BBL estimators exhibit substantial bias in estimating the investment cost parameter θ_x, with mean estimates of 0.633 (additive) and 0.542 (multiplicative) against a true value of 1.0. All three implementations of the BBL estimator

[20] In a previous version we followed Bajari et al. (2007) and Pakes et al. (2007) in first estimating exit and cost parameters, and then estimating entry parameters holding exit and cost parameters at their estimated values. The results were not meaningfully different. We report results from the one-step BBL estimator as it appears more directly comparable to our approach of solving (31).

[21] Increasing further the number of forward simulations and/or the number of periods did not appear to improve the accuracy of value function approximations.

[22] The distribution of the recursive estimator across Monte Carlo replications appears approximately Gaussian; see Appendix E for histograms.

Table 2: Estimation Results – BBL Model

Parameter	Value	Recursive	Additive	Multiplicative	Asymptotic
θ_x	1.000	1.001	0.633	0.542	-4.214
		(0.009)	(0.026)	(0.058)	(2.369)
F_ρ Lower Bound	22.000	21.756	20.372	18.848	22.397
		(0.439)	(1.157)	(1.036)	(3.999)
F_ρ Upper Bound	23.000	22.967	23.614	22.319	27.185
		(0.048)	(0.126)	(0.110)	(3.287)
F_κ Lower Bound	22.000	21.487	22.373	22.090	779.514
		(0.746)	(1.639)	(1.195)	(3504.104)
F_κ Upper Bound	30.000	32.314	33.742	27.087	1196.404
		(3.788)	(136.493)	(13.419)	(4953.821)

Parameter estimates for the BBL model. The "Parameter" column describes the parameters being estimated. θ_x parameterizes investment costs $c(x) = \theta_x x$. F_ρ and F_κ denote, respectively, the scrap value and entry cost distributions. The "Value" column reports the true value of those parameters in the data-generating process. The remaining columns report the mean and standard deviation of parameter estimates obtained from different estimators across 500 Monte Carlo simulations. The "Recursive" estimator is the nonlinear least squares estimator defined in equation (31). "Additive", "Multiplicative", and "Asymptotic" refer, respectively, to the BBL estimator with additive, multiplicative, and asymptotic deviations, as described in section 4.1.1.

perform reasonably well in estimating the scrap value parameters, with the implementation with asymptotic deviations performing the worst. The implementation of BBL with asymptotic deviations struggles to recover the entry cost parameters, whereas the additive implementation performs reasonably well in estimating the lower bound of F_κ but exhibits large variance in estimating the upper bound of F_κ. Out of the three BBL implementations, the multiplicative one performs the best in estimating F_κ. Crucially, Table 2 shows that the NLLS estimator performs best out of the four estimators we consider. Its finite sample bias is the smallest for all parameters, with the possible exception of the lower bound of F_κ, where the additive BBL estimator has a very similar bias. The standard deviations of the NLLS estimator are substantially smaller than those of the BBL estimators for all parameters, roughly by a factor of three to four.

Table 2 contains at least two noteworthy findings. First, the performance of the BBL estimator depends meaningfully on the choice of deviation policies. This is a drawback, as it may not be clear which class of deviations is most appropriate in a given application. Second, the NLLS estimator, which does away with the need to define deviation policies, outperforms all three BBL implementations.

4.2 An Extension of Hashmi and van Biesebroeck (2016)

In this section we extend the Hashmi and van Biesebroeck (2016) model of R&D in the automobile industry to allow for entry and exit. Importantly, and in contrast to the model of the preceding section, this model features shocks to the marginal cost of investment.

4.2.1 Static Price Competition

Suppose there are N single-product firms active in the market, indexed by $j = 1, \ldots, N$. Consumer i derives conditional indirect utility u_{ij} from purchasing firm j's product, where

$$u_{ij} = \begin{cases} \epsilon_i^{\text{out}} + (1-\varsigma)\varepsilon_{i0} & \text{if } j = 0 \\ \alpha p_j + \xi_j + \epsilon_i^{\text{in}} + (1-\varsigma)\varepsilon_{ij} & \text{if } j = 1, \ldots, N \end{cases} . \tag{32}$$

In (32), $j = 0$ denotes the outside good and p_j denotes good j's price. Goods are grouped into two nests, one containing all inside goods (i.e., those produced by one of the N firms) and one containing the outside good. The ε_{ij}'s are independent and identically distributed Type 1 Extreme Value random variables. The nest-level disturbances ϵ_i^{out} and ϵ_i^{in} follow the unique distribution such that $\epsilon_i^{\text{g}} + (1-\varsigma)\varepsilon_{ij}$, for g $= \{$in, out$\}$, is also Type 1 Extreme Value distributed – see Cardell (1997). The parameter $\varsigma \in [0, 1)$ is the nesting parameter. This is a standard nested-logit demand specification, leading to a well-known functional form for market shares.

A firm selling a good of quality ξ_j has constant marginal cost

$$mc(\xi_j) = \exp(\theta_{c1} + \theta_{c2}\xi_j) .$$

Firms compete à la Bertrand. There is a unique equilibrium to the pricing game (Caplin & Nalebuff, 1991), so that profits $\pi(\boldsymbol{\xi})$ as a function of firms' product qualities are well-defined.

4.2.2 Quality Transitions and Investment Decision

As before, firms' current investment affects their product quality in the following period. Product qualities are elements of $\Xi = \{-\infty, \xi_m, \xi_m + \delta, \ldots, \xi_M - \delta, \xi_M\}$. Transitions are as in the BBL model of Section 4.1, except that the probability of a positive quality shock is now allowed to depend not only on investment but also on current quality. Specifically, the probabiliity of a positive quality shock

$$u(\xi, x) = 1 - (1 + x)^{-\lambda(\xi)} , \tag{33}$$

where $\lambda(\xi) = \exp(\theta_{t2} + \theta_{t3}\xi + \theta_{t4}\xi^2)$. With parameters θ_t such that $\lambda(\xi)$ is a decreasing function of ξ over Ξ, this specification implies that the probability of quality improvement is decreasing in current quality, for any given level of investment. This captures the notion that it is harder to improve on a higher-quality product.

Firms face investment cost

$$c(x, \nu) = \theta_{x1}x + \theta_{x2}x^2 + \theta_{x3}x\nu , \tag{34}$$

and ν is drawn iid from a distribution F_ν known to the econometrician. The cost shock affects the marginal cost of investment and therefore affects the optimal investment choice. This rationalizes that two firms facing the same quality vector $\boldsymbol{\xi}$ may optimally choose different levels of investment. At the beginning of the period, prior to observing the cost shock ν, potential entrants privately observe their entry cost $\kappa \sim F_\kappa$, incumbents privately observe their scrap value $\rho \sim F_\rho$, and all firms make entry and exit decisions simultaneously. We assume that F_ρ and F_κ are both exponential distributions. Firms' entry and exit decisions determine the endogenous number of active firms $N \leq \bar{N}$, where $\bar{N}$ is the maximum number of firms in the market.

Other than including entry and exit, this model differs from the one in Hashmi and van Biesebroeck (2016) in two ways. First, our specification of upgrade probabilities in (33) differs from that in Hashmi and van Biesebroeck (2016). This is because their specification is not globally concave in investment, and thus does not meet the sufficient conditions for global concavity of the investment choice problem in Proposition 1 or Proposition 3. Second, our investment cost specification in (34) omits a cubic term present in Hashmi and van Biesebroeck (2016).

4.2.3 Parameterization

Table 3 reports parameters governing the data generating process. We fix $\bar{N} = 5$ in each simulated market, in line with simulations in Hashmi and van Biese-

Table 3: HvB Model Parameters

Parameter	**Value**
Data Structure	
Number of Quality Levels	15
Minimum Quality Level	-1.4
Maximum Quality Level	1.4
Distance Across Quality Levels (δ)	0.2
Maximum Number of Players	5
Number of Markets	100
Number of Periods	40
Model Parameters	
Discount Factor (β)	0.95
Demand System (γ_0, α)	[1.0, -0.222]
Marginal Cost (θ_c)	[2.47, 0.0]
Quality Transition (θ_t)	[0.347, -0.75, -0.3, -0.1]
Investment Cost (θ_x)	[2.625, 1.624, 0.5096]
Investment Cost Shock Distribution (F_ν)	Normal(0.0, 1.0)
Scrap Value Distribution (F_ρ)	Exponential(0.8)
Entry Costs Distribution (F_κ)	Exponential(11.0)
Forward Simulation Settings	
Number of Simulated Paths	250
Simulation Horizon	150

Parameterization of the Hashmi and van Biesebroeck (2016) model used in Monte Carlo simulations.

broeck (2016).[23] As in their model, incumbent product quality can take on fifteen values, from -1.4 to 1.4 in increments of 0.2. Firms' marginal cost of production is constant and equal to $mc_j = \exp(2.47)$, i.e., we set $\theta_{c1} = 2.47$ and $\theta_{c2} = 0$. The probability of a quality downgrade shock is $\theta_{t1} = 0.347$, and the remaining transition parameters are such that the probability of a quality upgrade shock is decreasing in own quality level at an increasing rate. Investment costs are convex ($\theta_{x2} > 0$) in investment and the marginal cost of investment is increasing in the shock ($\theta_{x3} > 0$). We follow Hashmi and van Biesebroeck (2016) in setting $F_\nu = N(0,1)$.[24] Our parameterization differs somewhat from that in Hashmi and van Biesebroeck (2016). For instance, we assume a lower probability of a quality downgrade shock. We also consider different values for investment cost parameters. For instance, HvB report a negative θ_{x2}. We impose a positive value to ensure convexity of the investment cost function. The changes in parameter values are not instrumental for the qualitative results we report in Section 4.2.4.

Finally, we must specify entry cost and scrap value distributions, which are not present in the Hashmi and van Biesebroeck (2016) model. We set $F_\rho = \text{Exp}(0.8)$ and $F_\kappa = \text{Exp}(11.0)$, where $\text{Exp}(\mu)$ denotes an exponential distribution with mean μ. The mean of the scrap value distribution is much smaller than that of entry costs. This is so that the DGP generates somewhat low entry and exit probabilities, as is commonly observed in empirical applications.[25]

4.2.4 Monte Carlo Results

We simulate 500 independent datasets from the model described in Section 4.2 and estimate the parameters of the model using the pseudo-MLE estimator defined in Section 3.1 and the implementations of the BBL estimator using additive, multiplicative, and asymptotic deviations, as discussed in Section

[23] In their empirical analysis, Hashmi and van Biesebroeck (2016) aggregate data to have a single market with 14 firms for the 1982-2006 period. When computing the equilibrium of the dynamic game, however, they restrict the number of firms to 5 to reduce computational burden. See Hashmi and van Biesebroeck (2016, Footnote 30).

[24] This implies that the marginal cost of investment can be negative, counter to Assumption 1. However, at the parameter values reported in Table 3, the probability that ν is sufficiently small to make the marginal cost of investment negative is less than 10^{-6}. We can straightforwardly accommodate different distributions.

[25] Table 3 details the distribution of scrap values and entry costs before they are scaled to have the same order of magnitude as firm profits. We implement this scaling by multiplying these shocks by the average of $\pi(\boldsymbol{\xi})/(1-\beta)$ across states.

Table 4: Parameter Estimates – HvB Model

Parameter	Value	Recursive	Additive	Multiplicative	Asymptotic
θ_{x1}	2.625	2.920	-4.972	-5.000	-4.829
		(0.261)	(0.360)	(0.000)	(1.069)
θ_{x2}	1.624	1.603	-5.000	-4.763	2.745
		(0.068)	(0.000)	(0.983)	(2.533)
θ_{x3}	0.510	0.548	9.324	8.733	0.232
		(0.023)	(2.424)	(3.259)	(1.420)
F_ρ Scale Parameter	0.800	0.754	0.901	0.702	0.721
		(0.061)	(0.031)	(0.079)	(0.035)
F_κ Scale Parameter	11.000	10.427	18.976	17.514	14.311
		(0.788)	(1.631)	(2.224)	(1.641)

Parameter estimates for the HvB model. The "Parameter" column describes the parameters being estimated. $\theta_{x1}, \theta_{x2}, \theta_{x3}$ parameterize investment costs: see (34). F_ρ and F_κ respectively denote scrap value and entry cost distributions. The "Value" column reports the true value of those parameters in the data-generating process. The remaining columns report the mean and standard deviation of parameter estimates obtained from different estimators across 500 Monte Carlo simulations. The "Recursive" estimator is the pseudo-MLE estimator described in Section 3.1. "Additive", "Multiplicative", and "Asymptotic" refer, respectively, to the BBL estimator with additive, multiplicative, and asymptotic deviations, as discussed in the main text.

4.1.1. We report the results in Table 4.[26] The results are now markedly different from those in Section 4.1. All implementations of the BBL estimator perform poorly in the estimation of the investment cost parameters, showing substantial finite sample bias and variance. The BBL estimator with asymptotic deviations performs best among the three BBL implementations, in contrast to the results in Section 4.1.2, but it still appears far from satisfactory. By contrast, the pseudo-MLE estimator based on recursive equilibrium conditions performs significantly better than the BBL estimators, showing much smaller finite sample bias and variance.[27]

The BBL estimators do a better job of estimating the scrap value and entry cost parameters. Again, the implementation with asymptotic deviations performs best. As in the case of investment cost parameters, the pseudo-MLE estimator significantly outperforms the three BBL estimators in the estimation of the entry cost and scrap value parameters.

[26]The distribution of the recursive estimator across Monte Carlo replications appears approximately Gaussian; see Appendix E for histograms.

[27]Appendix E.2 investigates the source of the remaining bias by replacing estimated first-stage objects with true MPE policies, finding that bias is largely attributable to first-stage estimation error.

5 Conclusion

We propose to estimate the parameters of dynamic games with continuous controls by exploiting the recursive optimality conditions characterizing firm behavior. Our focus is on models with shocks to the marginal cost of setting the continuous controls, though we also consider a model without such shocks. For the former class of models, we derive a pseudo maximum likelihood estimator and establish an existence result for symmetric Markov Perfect Equilibria. For the latter class of models, we use recursive equilibrium conditions to construct a nonlinear least squares estimator.

We then evaluate the performance of the NLLS and pseudo-MLE estimators in two Monte Carlo exercises based on Bajari et al. (2007) and Hashmi and van Biesebroeck (2016), respectively. In both exercises, we find that the estimators based on recursive optimality conditions significantly outperform three alternative implementations of the Bajari et al. (2007) estimator. In particular, we find that the BBL estimator has significant finite sample bias in estimating investment cost parameters, with estimates sometimes concentrated at the boundary of the parameter space. The estimators we propose, in contrast, exhibit substantially smaller finite sample bias and variance.

We have established these results in the context of models with a single continuous control, and under restrictions on the transitions of state variables. Though many models of interest in empirical industrial organization satisfy these restrictions, some do not. Extending our results to models with multiple continuous controls and more general transition processes is a fruitful avenue for future research.

References

Aguirregabiria, V., & Mira, P. (2007). Sequential estimation of dynamic discrete games. *Econometrica*, *75*, 1-53.

Bajari, P., Benkard, C. L., & Levin, J. (2007). Estimating dynamic models of imperfect competition. *Econometrica*, *75*, 1331-1370.

Berge, C. (1963). *Topological spaces: Including a treatment of multi-valued functions, vector spaces and convexity*. Oliver & Boyd.

Besanko, D., & Doraszelski, U. (2004). Capacity dynamics and endogenous asymmetries in firm size. *RAND Journal of Economics*, 23–49.

Caplin, A., & Nalebuff, B. (1991). Aggregation and imperfect competition: On the existence of equilibrium. *Econometrica*, *59*, 25.

Cardell, N. S. (1997). Variance components structures for the extreme-value and logistic distributions with application to models of heterogeneity. *Econometric Theory*, *13*(2), 185–213.

Çınlar, E. (2011). *Probability and stochastics*. Springer.

Denardo, E. V. (1967). Contraction mappings in the theory underlying dynamic programming. *SIAM Review*, *9*(2), 165–177.

Doraszelski, U., & Pakes, A. (2007). A framework for applied dynamic analysis in io. In M. P. R. Armstrong (Ed.), .

Doraszelski, U., & Satterthwaite, M. (2010). Computable markov-perfect industry dynamics. *The RAND Journal of Economics*, *41*(2), 215–243.

Egan, M., Hortaçsu, A., Kaplan, N., Sunderam, A., & Yao, V. (2025). Dynamic competition for sleepy deposits. *NBER Working Paper*(w34267).

Ericson, R., & Pakes, A. (1995). Markov-perfect industry dynamics: A framework for empirical work. *The Review of Economic Studies*, *62*, 53-82.

Fowlie, M., Reguant, M., & Ryan, S. P. (2016). Market-based emissions regulation and industry dynamics. *Journal of Political Economy*.

Hashmi, A. R., & van Biesebroeck, J. (2016). The relationship between market structure and innovation in industry equilibrium: A case study of the global automobile industry. *Review of Economics and Statistics*, *98*, 192-208.

Jofre-Bonet, M., & Pesendorfer, M. (2003). Estimation of a dynamic auction game. *Econometrica*, *71*(5), 1443–1489.

Klenke, A. (2013). *Probability theory: a comprehensive course* (2nd ed.). Springer.

Lim, C. S. H., & Yurukoglu, A. (2018). Dynamic Natural Monopoly Regulation: Time Inconsistency, Moral Hazard, and Political Environments. *Journal of Political Economy*, *126*(1). Retrieved 2018-12-11, from http://

www.journals.uchicago.edu/t-and-c
Liu, A. H., & Siebert, R. B. (2022, 1). The competitive effects of declining entry costs over time: Evidence from the static random access memory market. *International Journal of Industrial Organization*, *80*.
Pakes, A., & McGuire, P. (1994). Computing markov-perfect nash equilibria : Numerical implications of a dynamic differentiated product model. *The RAND Journal of Economics*, *25*, 555-589.
Pakes, A., Ostrovsky, M., & Berry, S. (2007). Simple estimators for the parameters of discrete dynamic games (with entry/exit examples). *RAND Journal of Economics*, *38*, 373-399.
Pesendorfer, M., & Schmidt-Dengler, P. (2008). Asymptotic least squares estimators for dynamic games. *The Review of Economic Studies*, *75*(3), 901–928.
Ryan, S. P. (2012). The costs of environmental regulation in a concentrated industry. *Econometrica*, *80*, 1019-1061.
Srisuma, S. (2013). Minimum distance estimators for dynamic games. *Quantitative Economics*, *4*(3), 549–583.
Stokey, N. L., Lucas, R. E., & Prescott, E. C. (1989). *Recursive methods in economic dynamics*. Harvard University Press.
Wang, Z., & Zhai, D. (2025). The Expansion and Dynamic Equilibrium Effects of Institutional Landlords.
Whitt, W. (1980). Representation and approximation of noncooperative sequential games. *SIAM Journal on Control and Optimization*, *18*(1), 33–48.

Appendices

Appendix A Proofs and Derivations

A.1 Simplifying the Incumbent's Continuation Value

This appendix derives equation (6), starting from (5), and establishes the technical conditions under which the derivation goes through. The argument proceeds in two steps.

Step 1: Conditional Independence. By Assumption CINS, the joint transition $F(\boldsymbol{\xi}' \mid \boldsymbol{\xi}, (1, x), \sigma_{-1}(\boldsymbol{\xi}, \boldsymbol{\varepsilon}_{-1}))$ factors into firm-specific transitions. This allows us to write

$$\int_{\boldsymbol{\xi}'} \bar{V}_I(\boldsymbol{\xi}') \, \mathrm{d}F(\boldsymbol{\xi}' \mid \boldsymbol{\xi}, (1, x), \sigma_{-1}(\boldsymbol{\xi}, \boldsymbol{\varepsilon}_{-1})) = \int_{\xi_1'} \widetilde{W}(\xi_1' \mid \boldsymbol{\xi}, \boldsymbol{\varepsilon}_{-1}, \sigma) \, \mathrm{d}F(\xi_1' \mid \xi_1, x) \;, \tag{35}$$

where

$$\widetilde{W}(\xi_1' \mid \boldsymbol{\xi}, \boldsymbol{\varepsilon}_{-1}, \sigma) := \int_{\xi_2'} \cdots \int_{\xi_{\bar{N}}'} \bar{V}_I(\xi_1', \boldsymbol{\xi}_{-1}') \, \mathrm{d}F(\xi_{\bar{N}}' \mid \xi_{\bar{N}}, \sigma(\boldsymbol{\xi}^{(\bar{N})}, \varepsilon_{\bar{N}})) \ldots \mathrm{d}F(\xi_2' \mid \xi_2, \sigma(\boldsymbol{\xi}^{(2)}, \varepsilon_2))$$

is the incumbent's ENPV given own future characteristic ξ_1' and the competitors' realized private shocks $\boldsymbol{\varepsilon}_{-1}$, when competitors behave according to σ.

Step 2: Independence of Private Shocks. Taking expectations with respect to $\boldsymbol{\varepsilon}_{-1}$ and using independence of the ε_i across firms, the $\boldsymbol{\varepsilon}_{-1}$ integral becomes a product of single-firm integrals. Changing the order of integration in (5), one encounters multiple terms of the form

$$\int_{\varepsilon_j} \int_{\xi_j'} \bar{V}_I(\boldsymbol{\xi}') \, \mathrm{d}F(\xi_j' \mid \xi_j, \sigma(\boldsymbol{\xi}^{(j)}, \varepsilon_j)) \, \mathrm{d}G_{\varepsilon_j} = \int_{\xi_j'} \bar{V}_I(\boldsymbol{\xi}') \, \mathrm{d}F^{\sigma}(\xi_j' \mid \boldsymbol{\xi}^{(j)}) \;, \tag{36}$$

where F^{σ} is the σ-induced marginal transition defined in (7). Substituting (36) into (35) and integrating $\widetilde{W}$ with respect to $\boldsymbol{\varepsilon}_{-1}$ replaces $\widetilde{W}$ with the W defined in (8), yielding equation (6).

Technical Conditions for (36). Equation (36) clearly holds when $F(\xi' \mid \xi, x)$ admits a density for all (ξ, x) or when Ξ is finite. It holds much more generally, via Fubini–Tonelli–type theorems for product measures constructed from a probability measure and a transition kernel, such as Theorem 6.3 in Chapter 1 of Çınlar (2011) or Theorem 14.29 of Klenke (2013). This implies that the discussion in the main text holds for continuous, discrete, and discrete-continuous public state transition processes.

For instance, the result in Çınlar (2011) is as follows.

Definition 3 (Transition kernel (Çınlar))**.** Let $(E, \mathcal{E})$ and $(F, \mathcal{F})$ be measurable spaces. A mapping $K : E \times \mathcal{F} \to \overline{\mathbb{R}}_+$ is called a *transition kernel* from $(E, \mathcal{E})$ into $(F, \mathcal{F})$ if

1. for every $B \in \mathcal{F}$, the mapping $x \mapsto K(x, B)$ is $\mathcal{E}$-measurable;
2. for every $x \in E$, the mapping $B \mapsto K(x, B)$ is a measure on $(F, \mathcal{F})$.

Theorem 1 (Measure–kernel–function (Çınlar, Thm. 6.3))**.** Let K be a transition kernel from $(E, \mathcal{E})$ into $(F, \mathcal{F})$. Then

$$Kf(x) := \int_F K(x, dy)\, f(y), \quad x \in E,$$

defines a function $Kf \in \mathcal{E}_+$ for every $f \in \mathcal{F}_+$. Moreover, for each measure μ on $(E, \mathcal{E})$,

$$\mu K(B) := \int_E \mu(dx)\, K(x, B), \quad B \in \mathcal{F},$$

defines a measure μK on $(F, \mathcal{F})$; and, for every measure μ on $(E, \mathcal{E})$ and $f \in \mathcal{F}_+$,

$$(\mu K)f \;=\; \mu(Kf) \;=\; \int_E \mu(dx) \int_F K(x, dy)\, f(y).$$

The last statement is precisely what we need. It says that the integral of f with respect to the mixture measure μK is equal to the mixture of integrals of f with respect to the kernels $K(x, \cdot)$, weighted by the measure μ. In the statement above, $\mathcal{E}_+$ and $\mathcal{F}_+$ denote the sets of non-negative measurable functions on $(E, \mathcal{E})$ and $(F, \mathcal{F})$, respectively. Our integrand in equation (36) is $\bar{V}_I(\xi_j, \boldsymbol{\xi}_{-j})$, which is non-negative in equilibrium. Theorem 14.29 of Klenke (2013) applies to integrable f.

A.2 Uniqueness of the Investment Decision: Proof of Proposition 1

We start with the proof of Proposition 1, restated below for convenience.

Proposition (1). If Assumption 1 holds and $F(\xi' \mid \xi, x)$ is UIC-admissible, then the problem

$$\max_{x\in[0,\bar{x}]} \pi(\boldsymbol{\xi}) - c(x,\nu) + \beta \int_{\xi_i'} W(\xi_i' \mid \boldsymbol{\xi}; F^\sigma)\, \mathrm{d}F(\xi_i' \mid \xi_i, x)$$

has a unique solution for all $\boldsymbol{\xi}$. Moreover, the maximizer is increasing in ν, and strictly so over the range of ν where it is interior.

Proof. Under UIC-admissibility, the linearity of the Riemann-Stieltjes integral with respect to the integrator implies

$$\begin{aligned}
\int_{\xi_i'} W(\xi_i' \mid \boldsymbol{\xi}; F^\sigma)\, \mathrm{d}F(\xi_i' \mid \xi_i, x) &= \int_{\xi_i'} W(\xi_i' \mid \boldsymbol{\xi}; F^\sigma)\, \mathrm{d}\left[L(\xi',\xi) + K(\xi',\xi)Q(\xi,x)\right] \\
&= \int_{\xi_i'} W(\xi_i' \mid \boldsymbol{\xi}; F^\sigma)\, \mathrm{d}L(\xi',\xi) \\
&\quad + Q(\xi,x) \underbrace{\int_{\xi_i'} W(\xi_i' \mid \boldsymbol{\xi}; F^\sigma)\, \mathrm{d}K(\xi',\xi)}_{A(\boldsymbol{\xi},F^\sigma)}
\end{aligned}$$

Let $A(\boldsymbol{\xi}, F^\sigma) := \int_{\xi_i'} W(\xi_i' \mid \boldsymbol{\xi}; F^\sigma)\, \mathrm{d}K(\xi',\xi)$. The investment first-order condition is then

$$-\frac{\partial c(x,\nu)}{\partial x} + \beta A(\boldsymbol{\xi}, F^\sigma) \frac{\partial Q(\xi,x)}{\partial x} = 0 \, ,$$

Suppose $A(\boldsymbol{\xi}, F^\sigma) \leq 0$. Then the objective function is strictly decreasing in x and the solution is $x^* = 0$. Suppose instead that $A(\boldsymbol{\xi}, F^\sigma) > 0$. Then the objective function is strictly concave in x:

$$-\frac{\partial^2 c(x,\nu)}{\partial x^2} + \beta A(\boldsymbol{\xi}, F^\sigma) \frac{\partial^2 Q(\xi,x)}{\partial x^2} < 0 \, .$$

Let $v(x, \boldsymbol{\xi}, \nu)$ denote the objective function. If $\partial_x v(0, \boldsymbol{\xi}, \nu) < 0$, then the solution is again $x^* = 0$. If $\partial_x v(\bar{x}, \boldsymbol{\xi}, \nu) > 0$, then $x^* = \bar{x}$. Otherwise, the solution is the unique solution to the first-order condition.

To establish the monotonicity statement, suppose first that $A(\boldsymbol{\xi}, F^\sigma) \leq 0$. Then $x^* = 0$ for all ν. Next suppose $A(\boldsymbol{\xi}, F^\sigma) > 0$. If $\partial_x v(0, \boldsymbol{\xi}, \nu_1) < 0$, and $\nu_1 < \nu_2$, then $\partial_x v(0, \boldsymbol{\xi}, \nu_2) < 0$ and in both cases $x^* = 0$ is optimal. If $\partial_x v(\bar{x}, \boldsymbol{\xi}, \nu_1) >$

0 and $\nu_1 < \nu_2$, then either $\partial_x v(\bar{x}, \boldsymbol{\xi}, \nu_2) \geq 0$, and $\bar{x}$ is optimal in both cases, or $\partial_x v(\bar{x}, \boldsymbol{\xi}, \nu_2) < 0$ and the solution is interior and thus less than $\bar{x}$. If $(\boldsymbol{\xi}, \nu)$ are such that the solution is interior, an application of the Implicit Function Theorem implies $\frac{\partial x^*(\boldsymbol{\xi},\nu)}{\partial \nu} < 0$. $\square$

We note here that under the assumption that the $W(\cdot)$ function is increasing, it is possible to give an alternative, and perhaps more intuitive, condition for strict concavity of the local income function. The $W(\cdot)$ function is an endogenous object, and imposing restrictions on it is a limitation of the following result. However, if the analyst is willing to make this monotonicity assumption – i.e., that the equilibrium is such that an increase in own characteristics is good in this precise sense – then the result below is useful for the estimation and solution of games whose transitions do not satisfy UIC-admissibility. Capacity accumulation games are a class of such games.

Assumption 3. Let Ξ be a compact subset of the real line, and denote its minimum and maximum by ξ_m and ξ_M. Let Ξ° be the interior of Ξ (understood in the discrete case as $\Xi \setminus \{\xi_m, \xi_M\}$). The family of distributions $F(\cdot \mid \xi, x)$ is such that, for all $\xi \in \Xi$ and $\xi' \in \Xi^\circ$,

(a) $F(\xi' \mid \xi, x)$ is twice-continuously differentiable in x.

(b) $F(\xi' \mid \xi, x)$ is strictly decreasing and strictly convex in x;

(c) $F(\xi_m \mid \xi, x)$ is decreasing and convex in x.

Assumption 3(a) is a technical condition required in the proof, but has little economic content. It is also nonrestrictive in practice, as in applications the econometrician typically imposes a parametric restriction on $F(\xi' \mid \xi, x)$ that satisfies this condition. Assumption 3(b) states that for any current and future characteristics (other than the endpoints of Ξ), an increase in investment x causes the cumulative distribution function to decrease; i.e., an increase in x increases a firm's distribution of future quality in the first-order stochastic dominance sense. Moreover, investment has decreasing marginal returns in the sense that the reduction in the CDF is decreasing in investment. Assumption 3(c) allows for the continuous case, where $F(\xi_m \mid \xi, x) = 0$ for all ξ and x, and for the discrete case in which $F(\xi_m \mid \xi, x)$ may be strictly positive and depend on x.

Assumption 3 is not strictly weaker than UIC-admissibility, but almost so in the following sense. A transition function could be UIC-admissible and fail to satisfy the monotonicity and curvature conditions in Assumption 3 if $K(\xi', \xi) > 0$. If transitions are so specified, investment increases the mass to the left of ξ' and must thus reduce the mass to its right. This would be a model where investment pushes the firm to specific points in the state space, rather than always enabling it to improve its state. This is less intuitive, we would argue, than a model where $K(\xi', \xi)$ is increasing in ξ': investment always moves mass to the right. Under that additional condition, UIC-admissibility implies Assumption 3.

Proposition 3. Suppose assumptions 1 and 3 hold. Moreover, suppose $W(\xi' \mid \boldsymbol{\xi}; F^\sigma)$ is increasing in ξ' for all $\boldsymbol{\xi}$. Then

$$v(x; \boldsymbol{\xi}, \nu) = \pi(\boldsymbol{\xi}) - c(x, \nu) + \beta \int_{\xi_1'} W(\xi_1' \mid \boldsymbol{\xi}; F^\sigma) \, \mathrm{d}F(\xi_1' \mid \xi_1, x)$$

is strictly concave in x.

Proof. By Integration by Parts,

$$\int_{\xi_1'} W(\xi_1' \mid \boldsymbol{\xi}; F^\sigma) \, \mathrm{d}F(\xi_1' \mid \xi_1, x) = -\int F(\xi' \mid \xi, x) \, \mathrm{d}W(\xi' \mid \boldsymbol{\xi}; F^\sigma) + W(\xi_M \mid \boldsymbol{\xi}; F^\sigma) \underbrace{F(\xi_M \mid \xi, x)}_{=1} - W(\xi_m \mid \boldsymbol{\xi}; F^\sigma) F(\xi_m \mid \xi, x)$$

Therefore,

$$\begin{aligned}\frac{\partial^2}{\partial x^2}\left(\int_{\xi_1'} W(\xi_1' \mid \boldsymbol{\xi}; F^\sigma) \, \mathrm{d}F(\xi_1' \mid \xi_1, x)\right) &= -\int \frac{\partial^2}{\partial x^2} F(\xi' \mid \xi, x) \, \mathrm{d}W(\xi' \mid \boldsymbol{\xi}; F^\sigma) \\ &\quad - W(\xi_m \mid \boldsymbol{\xi}; F^\sigma) \frac{\partial^2}{\partial x^2} F(\xi_m \mid \xi, x) \\ &< 0\end{aligned}$$

The differentiation under the integral sign is valid due to the twice continuous differentiability of $F(\xi' \mid \xi, x)$ in x. The inequality is due to Assumption 3 and the monotonicity of W. This inequality, coupled with the strict convexity of $c(x, \nu)$, implies the desired property. □

In addition to Assumption 3, we impose the standard restrictions on the investment cost function summarized in Assumption 1 and assume that $W(\xi' \mid \boldsymbol{\xi}; F^\sigma)$ is increasing. This says that the ENPV of starting the following period at ξ' conditional on $\boldsymbol{\xi}$ is increasing in ξ': that is, firms would rather start the following period from a higher rather than lower ξ, for all current states $\boldsymbol{\xi}$. This is a restriction on an equilibrium object. In theory, it may be violated. It could happen, for instance, that a higher quality in $t+1$ induces higher investment by competitors, and so much so that this reduces a firm's ENPV of profits.[28] That being said, for the models considered in this paper, we have not computed any equilibrium that violates this monotonicity condition.

A.3 Proof of Proposition 2

Consider an incumbent's Bellman equation after observing its scrap value:

$$V_I(\boldsymbol{\xi}, \rho) = \max\left\{\pi(\boldsymbol{\xi}) + \rho, \int \max_{x \in [0,\bar{x}]} h(x, \boldsymbol{\xi}, \nu, \bar{V}_I; F^\sigma)\, \mathrm{d}F_\nu\right\} , \tag{37}$$

where

$$\begin{aligned} h(x, \boldsymbol{\xi}, \nu, \bar{V}_I; F^\sigma) &:= \pi(\boldsymbol{\xi}) - c(x, \nu) + \\ &+ \beta \int_{\xi_1'} W(\xi_1' \mid \bar{V}_I, \boldsymbol{\xi}; F^\sigma)\, \mathrm{d}F(\xi_1' \mid \xi_1, x) , \end{aligned}$$

and

$$W(\xi_1' \mid \bar{V}_I, \boldsymbol{\xi}; F^\sigma) = \int_{\xi_2'} \cdots \int_{\xi_{\bar{N}}'} \bar{V}_I(\xi_1', \boldsymbol{\xi}_{-1}')\, \mathrm{d}F^\sigma(\xi_{\bar{N}}' \mid \boldsymbol{\xi}^{(\bar{N})}) \ldots \mathrm{d}F^\sigma(\xi_2' \mid \boldsymbol{\xi}^{(2)}) . \tag{8 – repeated}$$

Here we have chosen to make the dependence of W on $\bar{V}_I$ explicit to make the arguments that follow clearer.

Integrating (37) with respect to ρ establishes that, in equilibrium, $\bar{V}_I$ must be a fixed point of the operator $T_{F^\sigma} : \mathbb{R}^{|\boldsymbol{\Xi}^R|} \to \mathbb{R}^{|\boldsymbol{\Xi}^R|}$ defined by

$$T_{F^\sigma}\bar{V}_I(\boldsymbol{\xi}) := \int \max\left\{\pi(\boldsymbol{\xi}) + \rho, \int \max_{x \in [0,\bar{x}]} h(x, \boldsymbol{\xi}, \nu, \bar{V}_I; F^\sigma)\, \mathrm{d}F_\nu\right\}\, \mathrm{d}F_\rho . \tag{38}$$

Lemma 1. T_{F^σ} is a contraction.

Proof. T_{F^σ} satisfies Blackwell's sufficient conditions for a contraction (e.g., Stokey, Lucas, and Prescott (1989, Theorem 3.3)). Indeed,

[28] For a model where competitor investment can increase in a firm's state, see Besanko and Doraszelski (2004).

1. *Monotonicity*. If $\bar{V}_I' \geq \bar{V}_I$, then

$$h(x^*(\boldsymbol{\xi}, \nu, \bar{V}_I; F^\sigma), \boldsymbol{\xi}, \nu, \bar{V}_I; F^\sigma) \leq h(x^*(\boldsymbol{\xi}, \nu, \bar{V}_I; F^\sigma), \boldsymbol{\xi}, \nu, \bar{V}_I'; F^\sigma)$$
$$\leq \max_{x \in [0,\bar{x}]} h(x, \boldsymbol{\xi}, \nu, \bar{V}_I'; F^\sigma)$$

where x^* is the solution to the inner maximization problem in Equation (38). This implies that $T_{F^\sigma}\bar{V}_I(\boldsymbol{\xi}) \leq T_{F^\sigma}\bar{V}_I'(\boldsymbol{\xi})$ for all $\boldsymbol{\xi}$.

2. *Discounting*. Note that

$$T_{F^\sigma}(\bar{V}_I + a)(\boldsymbol{\xi}) = \int \max\left\{\pi(\boldsymbol{\xi}) + \rho, \int \max_{x \in [0,\bar{x}]} h(x, \boldsymbol{\xi}, \nu, \bar{V}_I; F^\sigma)\, \mathrm{d}F_\nu + \beta a\right\} \mathrm{d}F_\rho$$
$$\leq T_{F^\sigma}(\bar{V}_I)(\boldsymbol{\xi}) + \beta a$$

The boundedness condition in Blackwell's Theorem is satisfied because the functions $\bar{V}_I$ are simply vectors in a Euclidean space. □

It follows from Lemma 1 and the Contraction Mapping Theorem that to each F^σ there corresponds a unique integrated incumbent value function $\bar{V}_I(F^\sigma)$. We will need to establish that the map $\bar{V}_I(F^\sigma)$ is continuous. The next result will be useful in that regard. Let $\mathcal{F}$ be the set of collections of conditional distributions of firm quality conditional on the industry state. The elements of $\mathcal{F}$ are vectors of the form

$$\left[F(\xi' \mid \boldsymbol{\xi}) : \xi' \in \Xi, \boldsymbol{\xi} \in \boldsymbol{\Xi}^R\right] ,$$

where $0 \leq F(\xi' \mid \boldsymbol{\xi})$ and $\sum_{\xi' \in \Xi} F(\xi' \mid \boldsymbol{\xi}) = 1$ for all $\xi \in \Xi$ and $\boldsymbol{\xi} \in \boldsymbol{\Xi}^R$. Let $\mathcal{V} \subseteq \mathbb{R}^{|\boldsymbol{\Xi}^R|}$ be the set of possible integrated incumbent value functions $\bar{V}_I$.

Lemma 2. The map $T : \mathcal{F} \times \mathcal{V} \to \mathcal{V}$ with values $T_{F^\sigma}\bar{V}_I$ is continuous in F^σ.

Proof. The terms $W(\xi_1' \mid \bar{V}_I, \boldsymbol{\xi}; F^\sigma)$ are continuous in F^σ.[29] Therefore, the function h is continuous in F^σ and x.[30] The Maximum Theorem then implies that $\max_{x \in [0,\bar{x}]} h(x, \boldsymbol{\xi}, \nu, \bar{V}_I; F^\sigma)$ is continuous in F^σ. It then follows from equation (38) that $T_{F^\sigma}\bar{V}_I(\boldsymbol{\xi})$ is continuous in F^σ and thus, since $|\boldsymbol{\Xi}^R| < \infty$, that $T_{F^\sigma}\bar{V}_I$ is continuous in F^σ. □

[29] This is clear in the finite case we are restricting ourselves to. It also holds in greater generality. To see this, note that in the general case one can define T_{F^σ} in the space of bounded continuous functions on $\boldsymbol{\Xi}^R$ with the sup norm. For a continuous $\bar{V}_I$, continuity of $W(\xi_1' \mid \bar{V}_I, \boldsymbol{\xi}; F^\sigma)$ in F^σ follows from Helly's Second Theorem.

[30] Again, continuity in x is clear in the finite case. In the general case, one can again refer to Helly's Second Theorem.

Denardo (1967) states the result below in terms of policy functions (δ in his notation) and their associated return functions v_δ, which are the expected net present value of payoffs when following policy δ. The result holds in our setting. Denardo (1967)'s proof goes through unchanged.

Lemma 3 (Theorem 1, Denardo (1967))**.** For all $F \in \mathcal{F}$ and $\bar{V}_I \in \mathcal{V}$,

$$\|\bar{V}_I(F^\sigma) - \bar{V}_I\| \leq \frac{1}{1-\beta}\|T_{F^\sigma}\bar{V}_I - \bar{V}_I\| \; .$$

Lemma 4. The map $\bar{V}_I : \mathcal{F} \to \mathcal{V}$ with values $\bar{V}_I(F^\sigma)$ is continuous.

Proof. This is very similar to Lemma 3.2 of Whitt (1980). Let $F^\sigma \in \mathcal{F}$ and $\{F^\sigma_n\}_{n\in\mathbb{N}}$ be a sequence in $\mathcal{F}$ such that $\lim_{n\to\infty} F^\sigma_n = F^\sigma$. Then

$$\begin{aligned}\|\bar{V}_I(F^\sigma_n) - \bar{V}_I(F^\sigma)\| &\leq \frac{1}{1-\beta}\|T_{F^\sigma_n}\bar{V}_I(F^\sigma) - \bar{V}_I(F^\sigma)\| \\ &= \frac{1}{1-\beta}\|T_{F^\sigma_n}\bar{V}_I(F^\sigma) - T_{F^\sigma}\bar{V}_I(F^\sigma)\| \to 0\end{aligned}$$

as $n \to \infty$. The inequality is due to Lemma 3. The equality is due to $\bar{V}_I(F^\sigma)$ being a fixed point of T_{F^σ}. The limit is due to Lemma 2. □

Now consider the problem

$$\max_{x\in[0,\bar{x}]} h(x, \boldsymbol{\xi}, \nu, \bar{V}_I(F^\sigma); F^\sigma) \; . \tag{39}$$

Because of Proposition 1, this problem has a unique solution for each $\boldsymbol{\xi}$ and ν. Collect those in the policy function $\sigma^x(\boldsymbol{\xi}, \nu; F^\sigma)$. The value of problem (39) gives the ENPV of profits conditional on being active when competitor transitions are given by F^σ, $V^A_I(\boldsymbol{\xi}, \nu; F^\sigma)$, as in equation (4). By Berge's Maximum Theorem and the uniqueness of the optimal investment, $\sigma^x(\boldsymbol{\xi}, \nu; F^\sigma)$ and $V^A_I(\boldsymbol{\xi}, \nu; F^\sigma)$ are continuous in F^σ. The same is true of $V^A_E(\boldsymbol{\xi}_{-1}, \nu; F^\sigma)$.

It follows that

$$\bar{V}^A_s(\boldsymbol{\xi}; F^\sigma) := \int V^A_s(\boldsymbol{\xi}, \nu; F^\sigma)\, \mathrm{d}F_\nu \; , \;\; s \in \{E, I\}$$

are also continuous in F^σ. Therefore, the transitions implied by F^σ, given by

$$F(\xi' \mid \boldsymbol{\xi}; F^\sigma) = \begin{cases} 1 - F_\rho(\bar{V}_I^A(\boldsymbol{\xi}; F^\sigma) - \pi(\boldsymbol{\xi})) & \text{if } \xi > -\infty \\ & \text{and } \xi' = -\infty \\ F_\rho(\bar{V}_I^A(\boldsymbol{\xi}; F^\sigma) - \pi(\boldsymbol{\xi})) \int F(\xi' \mid \xi, \sigma^x(\boldsymbol{\xi}, \nu; F^\sigma)) \, \mathrm{d}F_\nu & \text{if } \xi > -\infty \\ & \xi' > -\infty \\ 1 - F_\kappa(\bar{V}_E^A(\boldsymbol{\xi}_{-1}; F^\sigma)) & \text{if } \xi = \xi' = -\infty \\ F_\kappa(\bar{V}_E^A(\boldsymbol{\xi}_{-1}; F^\sigma)) \int F(\xi' \mid \xi_e, \sigma^x(\boldsymbol{\xi}_{-1}, \nu; F^\sigma)) \, \mathrm{d}F_\nu & \text{if } \xi = -\infty \\ & \xi' > -\infty \end{cases} \tag{40}$$

are continuous in F^σ for all ξ' and $\boldsymbol{\xi}$.

We have thus established that the map $H : \mathcal{F} \to \mathcal{F}$ defined by (40) is continuous. It is clear that $\mathcal{F}$ is a compact and convex subset of $\mathbb{R}^{|\Xi| \times |\boldsymbol{\Xi}^R|}$. Therefore, Brouwer's Fixed-Point Theorem applies: there exists a $\bar{F}^\sigma \in \mathcal{F}$ such that $H(\bar{F}^\sigma) = \bar{F}^\sigma$.

To conclude the argument, consider the policy functions that obtain under $\bar{V}_I(\bar{F}^\sigma)$: the investment policy $\sigma^x(\boldsymbol{\xi}, \nu; \bar{F}^\sigma)$, the exit policy $\alpha^I(\boldsymbol{\xi}, \rho; \bar{F}^\sigma)$ that solves the outer maximization problem in (37), and the entry policy $\alpha^E(\boldsymbol{\xi}_{-1}, \kappa; \bar{F}^\sigma)$ that solves problem (14). By definition, they are optimal given $\bar{F}^\sigma$. Moreover, because $\bar{F}^\sigma$ is a fixed point of H, $\sigma^x(\boldsymbol{\xi}, \nu; \bar{F}^\sigma)$, $\alpha^I(\boldsymbol{\xi}, \rho; \bar{F}^\sigma)$, and $\alpha^E(\boldsymbol{\xi}_{-1}, \kappa; \bar{F}^\sigma)$ give rise to $\bar{F}^\sigma$. It follows that these policies are optimal against themselves, i.e., the strategy profile in which each firm plays these strategies is a symmetric Markov Perfect Equilibrium.[31]

A.4 Characterizing EV

This section derives equation (23).[32] We start from $V_I(\boldsymbol{\xi}, \rho)$:

$$V_I(\boldsymbol{\xi}, \rho) = \max \left\{ \pi(\boldsymbol{\xi}) + \rho, \bar{V}_I^A(\boldsymbol{\xi}) \right\} = \max_{\chi \in \{0,1\}} \left\{ \chi \bar{V}_I^A(\boldsymbol{\xi}) + (1 - \chi)(\pi(\boldsymbol{\xi}) + \rho) \right\} \tag{41}$$

[31] Strictly speaking, we construct a candidate equilibrium on a reduced state-space. One still has to check that the strategies induced on the original state space constitute an equilibrium of the game. See Doraszelski and Satterthwaite (2010), Section 6, in particular p. 239 (though the preceding pages give required preliminaries).

[32] Related calculations appear e.g. in Jofre-Bonet and Pesendorfer (2003); Pakes et al. (2007).

Letting $\alpha^I(\boldsymbol{\xi}, \rho)$ denote the optimal policy, we have

$$V_I(\boldsymbol{\xi}, \rho) = \alpha^I(\boldsymbol{\xi}, \rho)\bar{V}_I^A(\boldsymbol{\xi}) + (1 - \alpha^I(\boldsymbol{\xi}, \rho))(\pi(\boldsymbol{\xi}) + \rho)$$

and, integrating over ρ,

$$\begin{aligned} \bar{V}_I(\boldsymbol{\xi}) &:= \int V_I(\boldsymbol{\xi}, \rho)\, \mathrm{d}F_\rho \\ &= \mathbb{P}_I^A(\boldsymbol{\xi})\bar{V}_I^A(\boldsymbol{\xi}) + [1 - \mathbb{P}_I^A(\boldsymbol{\xi})]\pi(\boldsymbol{\xi}) \\ &+ [1 - \mathbb{P}_I^A(\boldsymbol{\xi})]\mathbb{E}[\rho \mid \rho > \bar{V}_I^A(\boldsymbol{\xi}) - \pi(\boldsymbol{\xi})] \end{aligned} \tag{42}$$

where $\mathbb{P}_I^A(\boldsymbol{\xi}) := \mathbb{P}(\alpha^I(\boldsymbol{\xi}, \rho) = 1)$ is the probability that an incumbent chooses to be active when its initial state is $\boldsymbol{\xi}$.

Next, letting $\sigma^x(\boldsymbol{\xi}, \nu)$ denote the optimal investment policy, we have that

$$V^A(\boldsymbol{\xi}, \nu) = \pi(\boldsymbol{\xi}) - c(\sigma^x(\boldsymbol{\xi}, \nu), \nu; \boldsymbol{\theta}_x) + \beta \sum_{\xi_1' \in \Xi} W(\xi_1' \mid \boldsymbol{\xi}, \boldsymbol{\theta}_{-\kappa}) P(\xi_1' \mid \xi_1, \sigma^x(\boldsymbol{\xi}, \nu)) \tag{43}$$

where, with slight abuse of notation, $P(\xi_1' \mid \xi, x)$ denotes the probability of the firm's own characteristic evolving from ξ_1 to ξ_1' when the firm invests x. Integrating (43) we get

$$\bar{V}_I^A(\boldsymbol{\xi}) = \pi(\boldsymbol{\xi}) - \int c(\sigma^x(\boldsymbol{\xi}, \nu), \nu; \boldsymbol{\theta}_x)\, \mathrm{d}F_\nu + \beta \sum_{\xi_1' \in \Xi} W(\xi_1' \mid \boldsymbol{\xi}, \boldsymbol{\theta}_{-\kappa}) \mathbb{P}^A(\xi_1' \mid \boldsymbol{\xi}) \tag{44}$$

where $\mathbb{P}^A(\xi_1' \mid \boldsymbol{\xi}) := \int P(\xi_1' \mid \xi_1, \sigma^x(\boldsymbol{\xi}, \nu))\, \mathrm{d}F_\nu$ is the ex-ante probability of ξ_1' given that the firm chooses to be active and invests optimally in state $\boldsymbol{\xi}$. Then, using the definition of $W(\xi_1' \mid \boldsymbol{\xi}, \boldsymbol{\theta}_{-\kappa})$, we have that

$$\begin{aligned} \bar{V}_I^A(\boldsymbol{\xi}) &= \pi(\boldsymbol{\xi}) - \int c(\sigma^x(\boldsymbol{\xi}, \nu), \nu; \boldsymbol{\theta}_x)\, \mathrm{d}F_\nu \\ &+ \beta \sum_{\xi_1' \in \Xi} \left(\sum_{\boldsymbol{\xi}_{-1}'} \bar{V}_I(\xi_1', \boldsymbol{\xi}_{-1}') \prod_{k>1} P^\sigma(\xi_k' \mid \boldsymbol{\xi}) \right) \mathbb{P}^A(\xi_1' \mid \boldsymbol{\xi})\,. \end{aligned} \tag{45}$$

where $P^\sigma(\xi_k' \mid \boldsymbol{\xi})$ is defined in equation (27).

We now plug (45) into (42) to obtain

$$
\begin{aligned}
\bar{V}_I(\boldsymbol{\xi}) = \pi(\boldsymbol{\xi}) - \mathbb{P}_I^A(\boldsymbol{\xi}) \int c(\sigma^x(\boldsymbol{\xi},\nu),\nu;\boldsymbol{\theta}_x)\,\mathrm{d}F_\nu \\
+ [1-\mathbb{P}_I^A(\boldsymbol{\xi})]\mathbb{E}[\rho \mid \rho > \bar{V}_I^A(\boldsymbol{\xi}) - \pi(\boldsymbol{\xi})] \\
+ \beta \sum_{\{\boldsymbol{\xi}':\xi_1'>-\infty\}} \bar{V}_I(\boldsymbol{\xi}') \prod_{k=1}^{\bar{N}} P^\sigma(\xi_k' \mid \boldsymbol{\xi})
\end{aligned}
\tag{46}
$$

where the last line uses the fact that $P^\sigma(\xi_1' \mid \boldsymbol{\xi}) = \mathbb{P}^A(\xi_1' \mid \boldsymbol{\xi})\mathbb{P}_I^A(\boldsymbol{\xi})$. Note that we can also make the sum in the last line over all $\boldsymbol{\xi}'$ if we define $\bar{V}_I(\boldsymbol{\xi}') = 0$ when $\xi_1' = -\infty$. With this convention it is more accurate to interpret $\bar{V}_I(\boldsymbol{\xi})$ as the expected net present value (ENPV) of landing in state $\boldsymbol{\xi}$ (given our assumption that firms that exit perish), rather than starting a period from state $\boldsymbol{\xi}$.

Observe that $\bar{\boldsymbol{V}}_I := [\bar{V}_I(\boldsymbol{\xi}) : \boldsymbol{\xi} \in \Xi^R, \xi_1 > -\infty]$ enters equation (46) in a non-linear fashion through the $\mathbb{E}[\rho \mid \rho > \bar{V}_I^A(\boldsymbol{\xi}) - \pi(\boldsymbol{\xi})]$ term. Fortunately, assuming that F_ρ is strictly increasing, we can deal with that term as follows:

$$
\begin{aligned}
\mathbb{E}[\rho \mid \rho > \bar{V}_I^A(\boldsymbol{\xi}) - \pi(\boldsymbol{\xi})] &= \mathbb{E}[\rho \mid F_\rho(\rho) > F_\rho(\bar{V}_I^A(\boldsymbol{\xi}) - \pi(\boldsymbol{\xi}))] \\
&= \mathbb{E}[\rho \mid F_\rho(\rho) > \mathbb{P}_I^A(\boldsymbol{\xi})] \\
&= \mathbb{E}[\rho \mid \rho > F_\rho^{-1}(\mathbb{P}_I^A(\boldsymbol{\xi}))] \,,
\end{aligned}
\tag{47}
$$

which does away with $\bar{V}_I^A(\boldsymbol{\xi})$. We can now plug (47) into (46) and stack across states with $\xi_1 > -\infty$:

$$
\bar{\boldsymbol{V}}_I = \boldsymbol{\pi} - \boldsymbol{K}(\boldsymbol{\theta}_x) + \boldsymbol{\Sigma}(F_\rho) + \beta \boldsymbol{M}(\boldsymbol{P})\bar{\boldsymbol{V}}_I \qquad \text{(23 – Repeated)}
$$

where the terms of this equation are defined in (24) to (27).

Finally, note that (23) does uniquely define $\bar{V}_I$ because the matrix $I - \beta\boldsymbol{M}(\boldsymbol{P})$ is invertible. Indeed, assume otherwise. Then there exists $x \in \Xi_I^R \setminus \{\mathbf{0}\}$ such that $[I-\beta\boldsymbol{M}(\boldsymbol{P})]x = 0$, or $x = \beta\boldsymbol{M}x$. This implies $\|x\|_\infty = \|\beta\boldsymbol{M}x\|_\infty$. However, letting $(\boldsymbol{M}x)_i$ denote the i-th coordinate of $\boldsymbol{M}x$, we have

$$
|(\boldsymbol{M}x)_i| = \left| \sum_{j=1}^{|\Xi_I^R|} M_{ij} x_j \right| \le \sum_{j=1}^{|\Xi_I^R|} M_{ij}|x_j| \le \|x\|_\infty \sum_{j=1}^{|\Xi_I^R|} M_{ij} \le \|x\|_\infty \,,
$$

where we have used that $\boldsymbol{M}$ is a sub-stochastic matrix, i.e. its rows sum to at most one.[33] Therefore $\|x\|_\infty = \|\beta\boldsymbol{M}x\|_\infty = \beta\|\boldsymbol{M}x\|_\infty \le \beta\|x\|_\infty$, a contradiction.

[33] The rows of $\boldsymbol{M}(\boldsymbol{P})$ need not sum to one because $\boldsymbol{M}(\boldsymbol{P})$ is the matrix of transitions between

Appendix B The Investment Contribution to the Likelihood

As shown in the main text, the conditional distribution of investment given $\boldsymbol{\xi}$ is given by

$$F_X(x \mid \boldsymbol{\xi}) = 1 - F_\nu\left((\sigma^x)^{-1}(x; \boldsymbol{\xi})\right) \ . \tag{19}$$

Here we build on the derivations in Appendix A.2. As shown there, the investment first-order condition can be written as

$$-\frac{\partial c(x, \nu; \boldsymbol{\theta}_x)}{\partial x} + \beta A(\boldsymbol{\xi}, F^\sigma; \boldsymbol{\theta}_{-\kappa})\frac{\partial Q(\xi, x)}{\partial x} = 0 \ ,$$

where $A(\boldsymbol{\xi}, F^\sigma; \boldsymbol{\theta}_{-\kappa}) = \int_{\xi'} W(\xi' \mid \boldsymbol{\xi}; F^\sigma, \boldsymbol{\theta}_{-\kappa}) \, \mathrm{d}K(\xi', \xi)$ and we have made explicit the dependence on the dynamic parameters to be estimated.

To characterize the likelihood at a parameter vector $\boldsymbol{\theta} = (\boldsymbol{\theta}_{-\kappa}, \boldsymbol{\theta}_\kappa)$, there are two cases to consider. First, suppose $A(\boldsymbol{\xi}, F^\sigma; \boldsymbol{\theta}_{-\kappa}) \leq 0$. Then, per the proof of Proposition 1, the optimal investment level is $x^* = 0$ for all ν. Therefore, the investment contribution to the likelihood is 1 if $x = 0$ and 0 otherwise.

Next, suppose $A(\boldsymbol{\xi}, F^\sigma; \boldsymbol{\theta}_{-\kappa}) > 0$. Then, again following the proof of Proposition 1, if $x = 0$ its contribution to the likelihood is given by equation (19). If $x > 0$, we need the conditional density. Differentiating (19) with respect to x gives the conditional density

$$f_X(x \mid \boldsymbol{\xi}) = -f_\nu\left((\sigma^x)^{-1}(x; \boldsymbol{\xi})\right) \cdot \frac{\partial}{\partial x}(\sigma^x)^{-1}(x; \boldsymbol{\xi}) \ . \tag{48}$$

The inverse $(\sigma^x)^{-1}(x; \boldsymbol{\xi})$ is characterized implicitly by the identity

$$\frac{\partial}{\partial x}c(x, (\sigma^x)^{-1}(x; \boldsymbol{\xi})) = \underbrace{\beta A(\boldsymbol{\xi}, F^\sigma; \boldsymbol{\theta}_{-\kappa})\frac{\partial Q(\xi, x)}{\partial x}}_{MB(\boldsymbol{\xi}, x)} \ . \tag{49}$$

Therefore,

$$(\sigma^x)^{-1}(x; \boldsymbol{\xi}) = (\partial_x c)^{-1}(MB(\boldsymbol{\xi}, x), x) \ , \tag{50}$$

where $(\partial_x c)^{-1}(\cdot, x)$ is the inverse of $\partial_x c(x, \nu)$ with respect to ν. We use (50) to compute the inverse policy. Moreover, applying the Implicit Function Theorem

states in $\Xi_I^R = \{\boldsymbol{\xi} \in \Xi^R : \xi_1 > -\infty\}$ rather than Ξ^R.

to (49) gives

$$\frac{\partial}{\partial x}(\sigma^x)^{-1}(x;\boldsymbol{\xi}) = \frac{-\partial_x^2 c(x,(\sigma^x)^{-1}(x;\boldsymbol{\xi})) + \partial_x MB(\boldsymbol{\xi},x)}{\partial_{\nu x}^2 c(x,(\sigma^x)^{-1}(x;\boldsymbol{\xi}))} \ .$$

As an illustration, let us apply the expression above to the specification in Section 4.2. We have

$$\begin{aligned} -\partial_x^2 c(x,(\sigma^x)^{-1}(x;\boldsymbol{\xi})) &= -2\theta_{x2} \\ \partial_x MB(\boldsymbol{\xi},x) &= \beta A(\boldsymbol{\xi})\partial_x^2 Q(\boldsymbol{\xi},x) \\ \partial_{\nu x}^2 c(x,(\sigma^x)^{-1}(x;\boldsymbol{\xi})) &= \theta_{x3} \ , \end{aligned}$$

so that

$$\frac{\partial}{\partial x}(\sigma^x)^{-1}(x;\boldsymbol{\xi}) = \frac{-2\theta_{x2} + \beta A(\boldsymbol{\xi})\partial_x^2 Q(\boldsymbol{\xi},x)}{\theta_{x3}} \ .$$

Appendix C First-Stage Estimation

The first stage produces estimates of the objects

$$\hat{\Phi} = \{\hat{\sigma}^x(\boldsymbol{\xi},\nu), \hat{\mathbb{P}}_I^A(\boldsymbol{\xi}), \hat{\mathbb{P}}_E^A(\boldsymbol{\xi}_{-1}), \hat{F}(\xi' \mid \xi, x), \hat{F}^\sigma(\xi' \mid \boldsymbol{\xi})\},$$

namely the investment policy function, the conditional probabilities that incumbents and entrants choose to be active, the quality transition kernel, and its policy-integrated counterpart. These objects feed the recursive estimator through the integrated value function in equation (23) and the BBL inequality estimator through the deviations described in Appendix D.2. This appendix discusses parameterization and estimation of first stage objects. Entry, exit, and transition primitives are estimated in the same way across the two designs. The investment policy is instead estimated by quantile regression in the HvB design and by linear regression in the BBL design. We do not report first-stage estimates because the estimators are standard and their role is exclusively to feed the second stage. Appendix E.2 reports a Monte Carlo exercise in which the estimated first stage is replaced by the true MPE policies, isolating the contribution of first-stage estimation error to second-stage performance.

Investment policy function. We follow Bajari et al. (2007) and exploit the identification result in equation (28), which relates the policy to the conditional quantile function of investment, $\sigma^x(\boldsymbol{\xi},\nu) = F_X^{-1}(1 - F_\nu(\nu) \mid \boldsymbol{\xi})$. In the HvB design, $\nu \sim N(0,1)$ has known continuous distribution, and we approximate the integral over ν in the integrated investment cost on a grid of Z quantile nodes $\{\nu_z, \omega_z\}_{z=1}^{Z}$ with $\omega_z = 1/Z$. At each $\tau_z = 1 - F_\nu(\nu_z)$ we run a quantile regression of investment on a vector of state features $f^x_{\text{HvB}}(\boldsymbol{\xi})$ that includes a constant, the firm's own quality (entered as a categorical variable with dummy coding), the number of active firms, the rank of the firm's quality, and the mean and maximum quality among rivals, and we set $\hat{\sigma}^x(\boldsymbol{\xi}, \nu_z) = \max\{0, f^x_{\text{HvB}}(\boldsymbol{\xi})'\hat{\chi}^x_{\tau_z}\}$. In the BBL design, since ν is degenerate at zero, we estimate the policy by linear regression on a richer feature vector $f^x_{\text{BBL}}(\boldsymbol{\xi})$ that adds bin-of-quality dummies and within-bin quadratics in own quality to the HvB regressors, setting $\hat{\sigma}^x(\boldsymbol{\xi}) = \max\{0, f^x_{\text{BBL}}(\boldsymbol{\xi})'\hat{\chi}^x\}$. In both instances, the non-negativity floor reflects the constraint $x \geq 0$ on investment.

Incumbent activity probability. We estimate the probability that an incumbent chooses to remain active by logistic regression, $\hat{\mathbb{P}}^A_I(\boldsymbol{\xi}) = \Lambda\big(f^I(\boldsymbol{\xi})'\hat{\chi}^I\big)$, where Λ is the logistic CDF and $f^I(\boldsymbol{\xi})$ contains a constant, the firm's own quality and its square, the number of active firms, the rank of the firm's quality, and the mean and maximum quality among rivals. The estimator and regressor specification are common to both designs.

Entrant activity probability. We estimate the probability that a potential entrant chooses to enter analogously, $\hat{\mathbb{P}}^A_E(\boldsymbol{\xi}_{-1}) = \Lambda\big(f^E(\boldsymbol{\xi}_{-1})'\hat{\chi}^E\big)$, with $f^E(\boldsymbol{\xi}_{-1})$ containing a constant, the number of active firms, and the rank, mean, and maximum of rival quality. The estimator and regressor specification are common to both designs.

Quality transition kernel. The state-conditional transition kernel $F(\xi' \mid \xi, x)$ governing how next-period quality depends on current quality and investment is parameterised by a low-dimensional vector of transition primitives $\boldsymbol{\theta}_t$. We estimate $\hat{\boldsymbol{\theta}}_t$ by maximum likelihood on observed quality transitions and obtain $\hat{F}(\xi' \mid \xi, x)$ by plugging $\hat{\boldsymbol{\theta}}_t$ into the parametric kernel. The two designs share this estimator but use different parameterisations of $F(\cdot \mid \xi, x)$: both designs use equation (30), but they feature different quality upgrade functions $u(\xi, x)$ as reported in the main text.

Policy-integrated transition kernel. The integrated value function depends on the kernel $F^\sigma(\xi' \mid \boldsymbol{\xi})$ that gives the distribution of a firm's next-period quality conditional on its current state, integrating over the firm's investment policy. We construct the estimate by integrating the estimated transition kernel against the estimated policy on the same quadrature grid $\{\nu_z, \omega_z\}$ used in the cost integral:

$$\hat{F}^{\sigma}(\xi' \mid \boldsymbol{\xi}) = \sum_z \omega_z \hat{F}\big(\xi' \mid \xi, \hat{\sigma}^x(\boldsymbol{\xi}, \nu_z)\big) \ .$$

In the BBL design the sum collapses to a single term because ν is degenerate. The construction is otherwise identical across the two designs.

Appendix D A Review of BBL and Implementation Details

D.1 A Review of the BBL Inequality Estimator

This subsection reviews the Bajari et al. (2007) estimator, the point of comparison for the estimators in Section 3.1. This estimator exploits the fact that, in an MPE, no deviation from equilibrium policy can improve a firm's expected discounted profits.

The expected discounted stream of profits of an incumbent playing strategy $\tilde{\sigma}^I = (\tilde{\sigma}^x, \tilde{\alpha}^I)$ when all its competitors play strategy $\sigma = (\sigma^x, \alpha^I, \alpha^E)$ is given by

$$\bar{V}_I(\boldsymbol{\xi}; \tilde{\sigma}^I, \sigma, \boldsymbol{\theta}_{-\kappa}) = \mathbb{E}\left\{ \sum_{t=0}^{\tau_e} \beta^t \left[\pi(\boldsymbol{\xi}_t) - c(\tilde{\sigma}^x(\boldsymbol{\xi}_t, \nu_t), \nu_t; \boldsymbol{\theta}_x)\right] + \beta^{\tau_e} \rho_{\tau_e} \,\middle|\, \boldsymbol{\xi}_0 = \boldsymbol{\xi} \right\} , \tag{51}$$

where $\tau_e := \min\{t : \tilde{\sigma}^I(\boldsymbol{\xi}_t, \rho_t) = 0\}$ is the incumbent's potentially infinite exit date and the public state evolves according to the probability distribution induced by the policy functions $(\tilde{\sigma}^I, \sigma)$ and F_ν, F_ρ, F_κ. The expected discounted stream of profits of a potential entrant playing strategy $\tilde{\sigma}^E = (\tilde{\sigma}^x, \tilde{\alpha}^I, \tilde{\alpha}^E)$ when all its competitors play strategy σ is

$$\bar{V}_E(\boldsymbol{\xi}; \tilde{\sigma}^E, \sigma, \boldsymbol{\theta}) = \int\int \tilde{\alpha}^E(\boldsymbol{\xi}, \kappa) v(\kappa, \nu, \boldsymbol{\xi}; \boldsymbol{\theta}_{-\kappa}) \, \mathrm{d}F_\nu \, \mathrm{d}F_\kappa(\boldsymbol{\theta}_\kappa) \ ,$$

where

$$\begin{aligned} v(\kappa, \nu, \boldsymbol{\xi}; \boldsymbol{\theta}_{-\kappa}) &:= -\kappa - c(\tilde{\sigma}^x(\boldsymbol{\xi}, \nu), \nu; \boldsymbol{\theta}_x) \\ &\quad + \beta \int\int \bar{V}_I(\boldsymbol{\xi}'; \tilde{\sigma}^E, \sigma, \boldsymbol{\theta}_{-\kappa}) \, \mathrm{d}F^{\sigma}(\boldsymbol{\xi}'_{-1} \mid \boldsymbol{\xi}) \, \mathrm{d}F(\xi'_1 \mid \xi_1, \tilde{\sigma}^x(\boldsymbol{\xi}, \nu)) \ . \end{aligned}$$

A symmetric strategy profile $(\sigma, \ldots, \sigma)$ is a Symmetric Markov Perfect Equilibrium only if, for all $\boldsymbol{\xi}$ and σ',

$$\bar{V}_I(\boldsymbol{\xi}; \sigma, \sigma, \boldsymbol{\theta}_{-\kappa}) \geq \bar{V}_I(\boldsymbol{\xi}; \sigma', \sigma, \boldsymbol{\theta}_{-\kappa}) \quad \text{and} \quad \bar{V}_E(\boldsymbol{\xi}; \sigma, \sigma, \boldsymbol{\theta}) \geq \bar{V}_E(\boldsymbol{\xi}; \sigma', \sigma, \boldsymbol{\theta}) \; . \tag{52}$$

Bajari et al. (2007) base their estimator on the equilibrium conditions (52).[34] Though their estimator can in principle be used to estimate set-identified models, it has not, to our knowledge, been applied as such. We thus focus on point-identified models, which we now define. Let $\mathcal{E}(\boldsymbol{\theta})$ be the set of SMPEs when the parameters of the model are given by $\boldsymbol{\theta}$.

Assumption 4 (Identification)**.** For any $\boldsymbol{\theta}, \boldsymbol{\theta}' \in \Theta$, $\mathcal{E}(\boldsymbol{\theta}) \cap \mathcal{E}(\boldsymbol{\theta}') = \emptyset$.

Given a state $\boldsymbol{\xi}$ and policy functions σ, σ', define

$$g(\boldsymbol{\xi}, \sigma', \sigma; \boldsymbol{\theta}) := \bar{V}(\boldsymbol{\xi}; \sigma, \sigma, \boldsymbol{\theta}) - \bar{V}(\boldsymbol{\xi}; \sigma', \sigma, \boldsymbol{\theta}) \; ,$$

where $\bar{V}$ ought to be interpreted as either $\bar{V}_I$ or $\bar{V}_E$ depending on the first coordinate of $\boldsymbol{\xi}$. Let H be a distribution over the space of pairs $(\boldsymbol{\xi}, \sigma')$. Define

$$Q(\boldsymbol{\theta}, \sigma) := \int \left(\min \left\{ g(\boldsymbol{\xi}, \sigma', \sigma; \boldsymbol{\theta}), 0 \right\} \right)^2 \mathrm{d}H(\boldsymbol{\xi}, \sigma') \; . \tag{53}$$

Let $\boldsymbol{\theta}_0$ denote the true parameters of the model. If $\sigma \in \mathcal{E}(\boldsymbol{\theta}_0)$, the equilibrium conditions above imply that $Q(\boldsymbol{\theta}_0, \sigma) = 0$. Under Assumption 4, $\sigma \in \mathcal{E}(\boldsymbol{\theta}_0) \Rightarrow \sigma \notin \mathcal{E}(\boldsymbol{\theta}')$ if $\boldsymbol{\theta}' \neq \boldsymbol{\theta}_0$. Therefore, if $\boldsymbol{\theta}' \neq \boldsymbol{\theta}_0$, then there must exist $(\boldsymbol{\xi}, \sigma')$ for which $g(\boldsymbol{\xi}, \sigma', \sigma; \boldsymbol{\theta}') < 0$. It follows that, for an appropriate choice of H, $Q(\boldsymbol{\theta}', \sigma) > 0$.[35]

Bajari et al. (2007) propose estimating the structural parameters of the model by minimizing a sample analog of (53). In particular, given a set of $\{(\boldsymbol{\xi}_i, \sigma'_i)\}_{i=1}^{n_I}$

[34] Condition (52) is slightly weaker than Markov Perfect Equilibrium as it allows violations of optimality at sets of measure zero (according to F_ρ and F_κ).

[35] To be more precise, $Q(\boldsymbol{\theta}', \sigma) > 0$ requires that $g(\boldsymbol{\xi}, \sigma', \sigma; \boldsymbol{\theta}') < 0$ on a set of positive H-measure for all $\boldsymbol{\theta}' \neq \boldsymbol{\theta}_0$. We can attach this condition to our definition of MPE. Given a measure μ on the set of tuples $(\boldsymbol{\xi}, \sigma')$, say that $(\sigma, \ldots, \sigma)$ is a symmetric MPE if $g(\boldsymbol{\xi}, \sigma', \sigma; \boldsymbol{\theta}_0) < 0$ with zero μ-measure. Then choose H such that μ is absolutely continuous with respect to H. If $\boldsymbol{\theta}' \neq \boldsymbol{\theta}_0$, Assumption 4 implies that $g(\boldsymbol{\xi}, \sigma', \sigma; \boldsymbol{\theta}') < 0$ with positive μ-measure. This implies that $g(\boldsymbol{\xi}, \sigma', \sigma; \boldsymbol{\theta}') < 0$ with positive H-measure, otherwise absolute continuity of μ with respect to H would be violated. Thus, $Q(\boldsymbol{\theta}', \sigma) > 0$. This hints at difficulties with the BBL approach: the measure H has to be rich, in the sense of $\mu \ll H$, where μ is itself rich enough that we are willing to define MPE on its basis. If H is not sufficiently rich, the equilibrium conditions may be violated at a set of positive μ-measure that is neglected by H. In this case $Q(\boldsymbol{\theta}', \sigma) = 0$.

pairs and an estimate of the strategy profile $\hat{\sigma}$, they propose minimizing

$$\hat{Q}(\boldsymbol{\theta}, \hat{\sigma}) := \frac{1}{n_I} \sum_{i=1}^{n_I} \left(\min \left\{ g(\boldsymbol{\xi}_i, \sigma_i', \hat{\sigma}; \boldsymbol{\theta}), 0 \right\} \right)^2 . \tag{54}$$

Evaluating this objective requires estimates of $\bar{V}(\boldsymbol{\xi}; \sigma', \hat{\sigma}, \boldsymbol{\theta})$. We review Bajari et al. (2007)'s proposal to obtain these estimates and alternatives in the next subsection.

D.2 BBL Implementation Details

We estimate the first-stage policy function $\hat{\sigma}^x(\boldsymbol{\xi}, \nu_\tau)$ and the entry and exit probabilities $\hat{\mathbb{P}}_E^A(\boldsymbol{\xi}_{-1}), \hat{\mathbb{P}}_I^A(\boldsymbol{\xi})$ as described in Appendix C, where $\tau \in (0,1)$ indexes the quantile of F_ν and $\hat{\chi}_\tau^x$, $\hat{\chi}^E$, $\hat{\chi}^I$ denote the corresponding regression coefficient estimates.

The BBL variants we consider differ in the construction of the deviations from these policies, which we denote as $\{\tilde{\sigma}^x(\boldsymbol{\xi}, \nu_\tau), \tilde{\mathbb{P}}_E^A(\boldsymbol{\xi}_{-1}), \tilde{\mathbb{P}}_I^A(\boldsymbol{\xi})\}$. Let i index a deviation. In what we term Additive BBL we draw $o_i^x \sim N(0, 0.3)$, $o_i^E \sim N(0, 0.5)$, and $o_i^I \sim N(0, 0.5)$ for each inequality i, as in Bajari et al. (2007). We then form deviations as

$$\begin{aligned}
\tilde{\sigma}_i^x(\boldsymbol{\xi}, \nu_\tau) &= \max\{0, \hat{\sigma}^x(\boldsymbol{\xi}, \nu_\tau) + o_i^x\} \\
\tilde{\mathbb{P}}_{E,i}^A(\boldsymbol{\xi}_{-1}) &= \min\{\max\{0, \hat{\mathbb{P}}_E^A(\boldsymbol{\xi}_{-1}) + o_i^E\}, 1\} \\
\tilde{\mathbb{P}}_{I,i}^A(\boldsymbol{\xi}) &= \min\{\max\{0, \hat{\mathbb{P}}_I^A(\boldsymbol{\xi}) + o_i^I\}, 1\} .
\end{aligned}$$

In what we term Multiplicative BBL we form deviations as

$$\begin{aligned}
\tilde{\sigma}_i^x(\boldsymbol{\xi}, \nu_\tau) &= \iota_i^x \hat{\sigma}^x(\boldsymbol{\xi}, \nu_\tau) \\
\tilde{\mathbb{P}}_{E,i}^A(\boldsymbol{\xi}_{-1}) &= \iota_i^E \hat{\mathbb{P}}_E^A(\boldsymbol{\xi}_{-1}) \\
\tilde{\mathbb{P}}_{I,i}^A(\boldsymbol{\xi}) &= \iota_i^I \hat{\mathbb{P}}_I^A(\boldsymbol{\xi}) ,
\end{aligned}$$

where $\iota_i^x, \iota_i^E, \iota_i^I$ are sampled uniformly and independently from the vector $\{0.90, 0.95, 1.05, 1.10\}$, as in Hashmi and van Biesebroeck (2016). Lastly, in what we term Asymptotic BBL we draw $\tilde{\boldsymbol{\chi}}_i \sim \mathrm{N}(\hat{\boldsymbol{\chi}}, \varsigma \hat{\boldsymbol{\Sigma}}_{\boldsymbol{\chi}})$ for each i, where ς is a scaling constant we set equal to 1. We then form deviations as

$$\begin{aligned}\tilde{\sigma}_i^x(\boldsymbol{\xi},\nu_\tau) &= f^x(\boldsymbol{\xi})'\tilde{\chi}_{i\tau}^x\\ \tilde{\mathbb{P}}_{E,i}^A(\boldsymbol{\xi}_{-1}) &= \Lambda\big(f^E(\boldsymbol{\xi}_{-1})'\tilde{\chi}_i^E\big)\\ \tilde{\mathbb{P}}_{I,i}^A(\boldsymbol{\xi}) &= \Lambda\big(f^I(\boldsymbol{\xi})'\tilde{\chi}_i^I\big).\end{aligned}$$

Note that all three sets of BBL deviations require the econometrician to choose tuning parameters: the distributions of additive deviations, the vector of constants constituting the multiplicative deviations, and the scaling constant for asymptotic deviations. The estimators we present in this paper dispense with this requirement.

Forward simulation of potentially non-linear value functions. Given estimated policy functions and deviations, typical implementations of the BBL inequality estimator compute value functions by forward simulation. As shown by equation (23), the incumbent value function depends on the conditional expectation of scrap values, which may cause the value function to be non-linear in model parameters. This nonlinearity increases the computational cost of the forward simulation routine, as the discussion in Section 3.2 no longer applies.

However, it is still the case that the forward simulation routine can be performed essentially once. To see this, note that we can define the exit policy as a function of a $U[0,1]$ random variable by means of a change of variable: $\check{\sigma}^I(\boldsymbol{\xi},\tau) := \sigma^I(\boldsymbol{\xi},F_\rho^{-1}(\tau))$. It follows that $\check{\sigma}^I(\boldsymbol{\xi},\tau) = \mathbb{1}\{\tau \leq \mathbb{P}_I^A(\boldsymbol{\xi})\}$.[36] We thus take $\tau \sim U[0,1]$ draws and use those to simulate exit decisions: incumbents remain active if and only if $\tau \leq \mathbb{P}_I^A(\boldsymbol{\xi})$. As we vary structural parameters, these simulations do not need to be repeated. All that needs to be recomputed is the scrap value that accrues to firms when they do decide to exit, which is $F_\rho^{-1}(\tau;\boldsymbol{\theta}_\rho)$, when $\tau > \mathbb{P}_I^A(\boldsymbol{\xi})$. The cost of repeatedly calling F_ρ^{-1} is typically dwarfed by the cost of repeating the simulation.

Appendix E Recursive Estimator Additional Results

This appendix presents additional Monte Carlo results for the recursive estimator. Section E.1 displays the full distribution of estimates across replications for both the BBL and HvB designs. Section E.2 isolates the contribution of first

[36] Indeed, $\check{\sigma}^I(\boldsymbol{\xi},\tau) = \sigma^I(\boldsymbol{\xi},F_\rho^{-1}(\tau)) = \mathbb{1}\{F_\rho^{-1}(\tau) \leq \bar{V}_I^A(\boldsymbol{\xi}) - \pi(\boldsymbol{\xi})\} = \mathbb{1}\{\tau \leq F_\rho(\bar{V}_I^A(\boldsymbol{\xi}) - \pi(\boldsymbol{\xi}))\} = \mathbb{1}\{\tau \leq \mathbb{P}_I^A(\boldsymbol{\xi})\}$.

stage estimation error by replacing estimated policy functions with true MPE policies.

E.1 Distribution of Estimates

Figures 1 and 2 display histograms of the recursive estimator across Monte Carlo replications for the BBL and HvB designs, respectively. The distributions are centered around the true parameter values in both designs. Some bias is visible for θ_{x3} in the HvB design, which may in part be attributable to first stage estimation error, consistent with the evidence in Section E.2.

E.2 Oracle First Stage

To isolate the contribution of first stage estimation error to second-stage bias, we conduct a Monte Carlo exercise using an "oracle" first stage that replaces the estimated policy functions with the true MPE policies. Table 5 and Figure 3 compare the pseudo-MLE estimates under the estimated first stage versus the oracle first stage. Under the oracle first stage, every mean estimate is within 0.6% of the true value (the largest absolute deviation, on θ_{x1}, is 0.013). Standard deviations also fall for four of the five parameters; the exception is θ_{x3}, where a tighter but biased estimated-first stage distribution gives way to a wider but unbiased one when first stage error is removed. The pattern indicates that the pseudo-MLE bias reported in the main results is attributable to first stage estimation error.

Table 5: Oracle vs Production First Stage – HvB Design

Parameter	Value	Estimated First Stage	Oracle First Stage
θ_{x1}	2.625	2.920	2.638
		(0.261)	(0.216)
θ_{x2}	1.624	1.603	1.621
		(0.068)	(0.046)
θ_{x3}	0.510	0.548	0.512
		(0.023)	(0.046)
F_ρ Scale Parameter	0.800	0.754	0.799
		(0.061)	(0.016)
F_κ Scale Parameter	11.000	10.427	10.991
		(0.788)	(0.304)

Mean estimates and standard deviations (in parentheses) across 500 Monte Carlo replications.

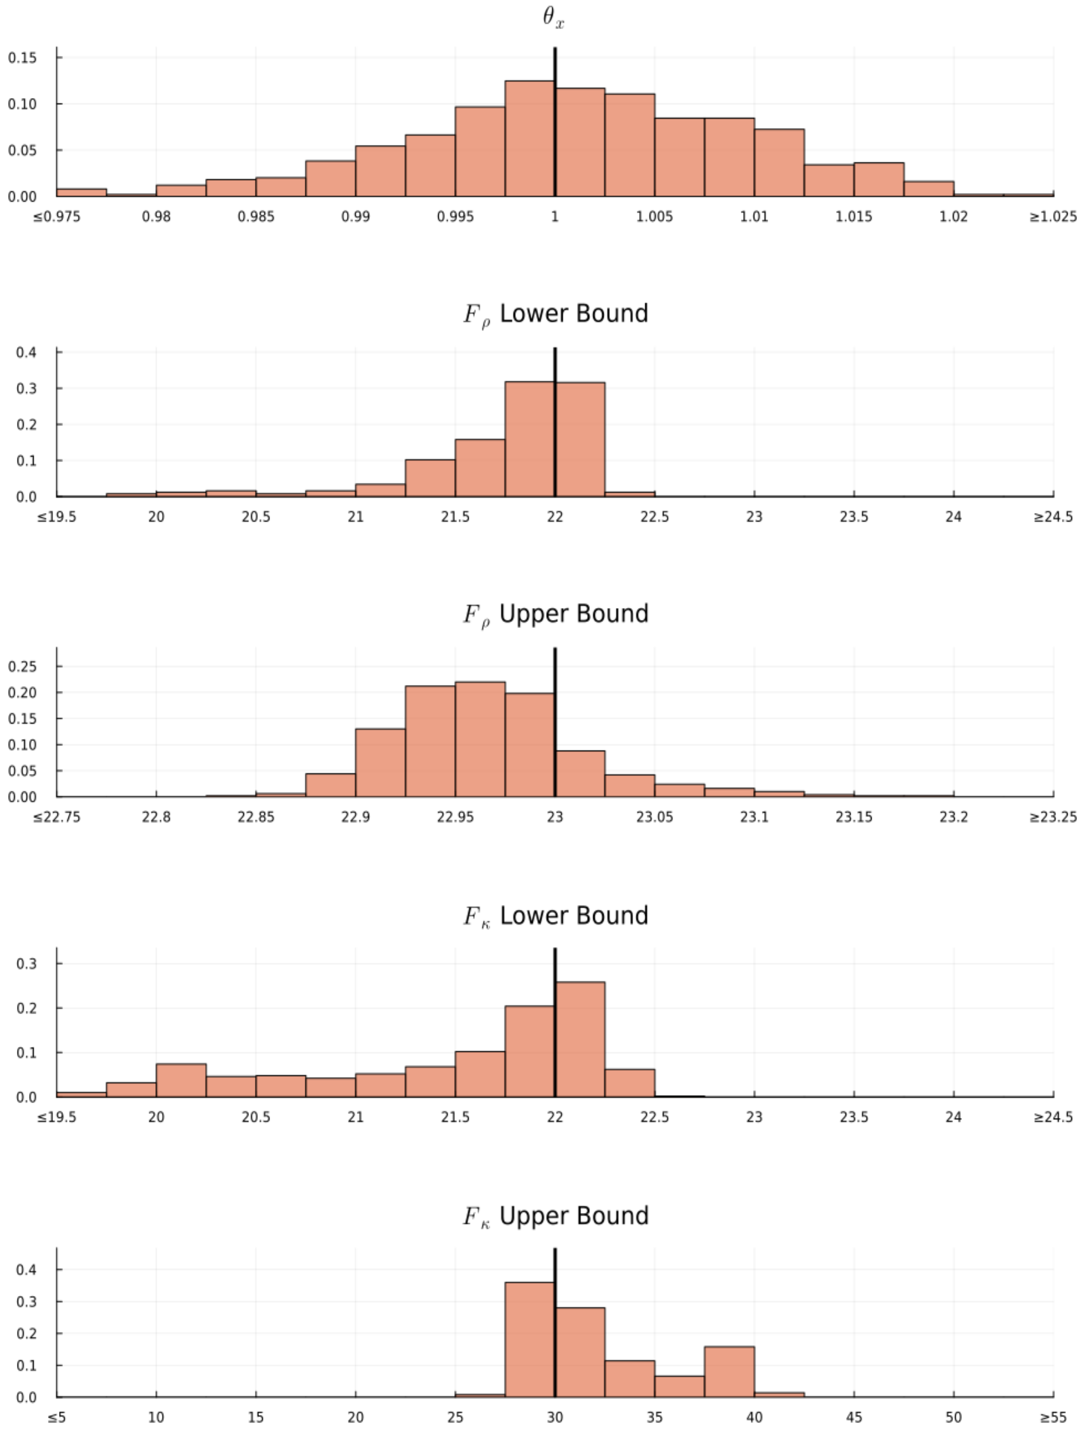

Figure 1: Recursive Estimator – BBL Design

Distribution of the NLLS recursive estimator across 500 Monte Carlo replications for the BBL design described in Section 4.1. The true parameter values are indicated by the vertical dashed lines.

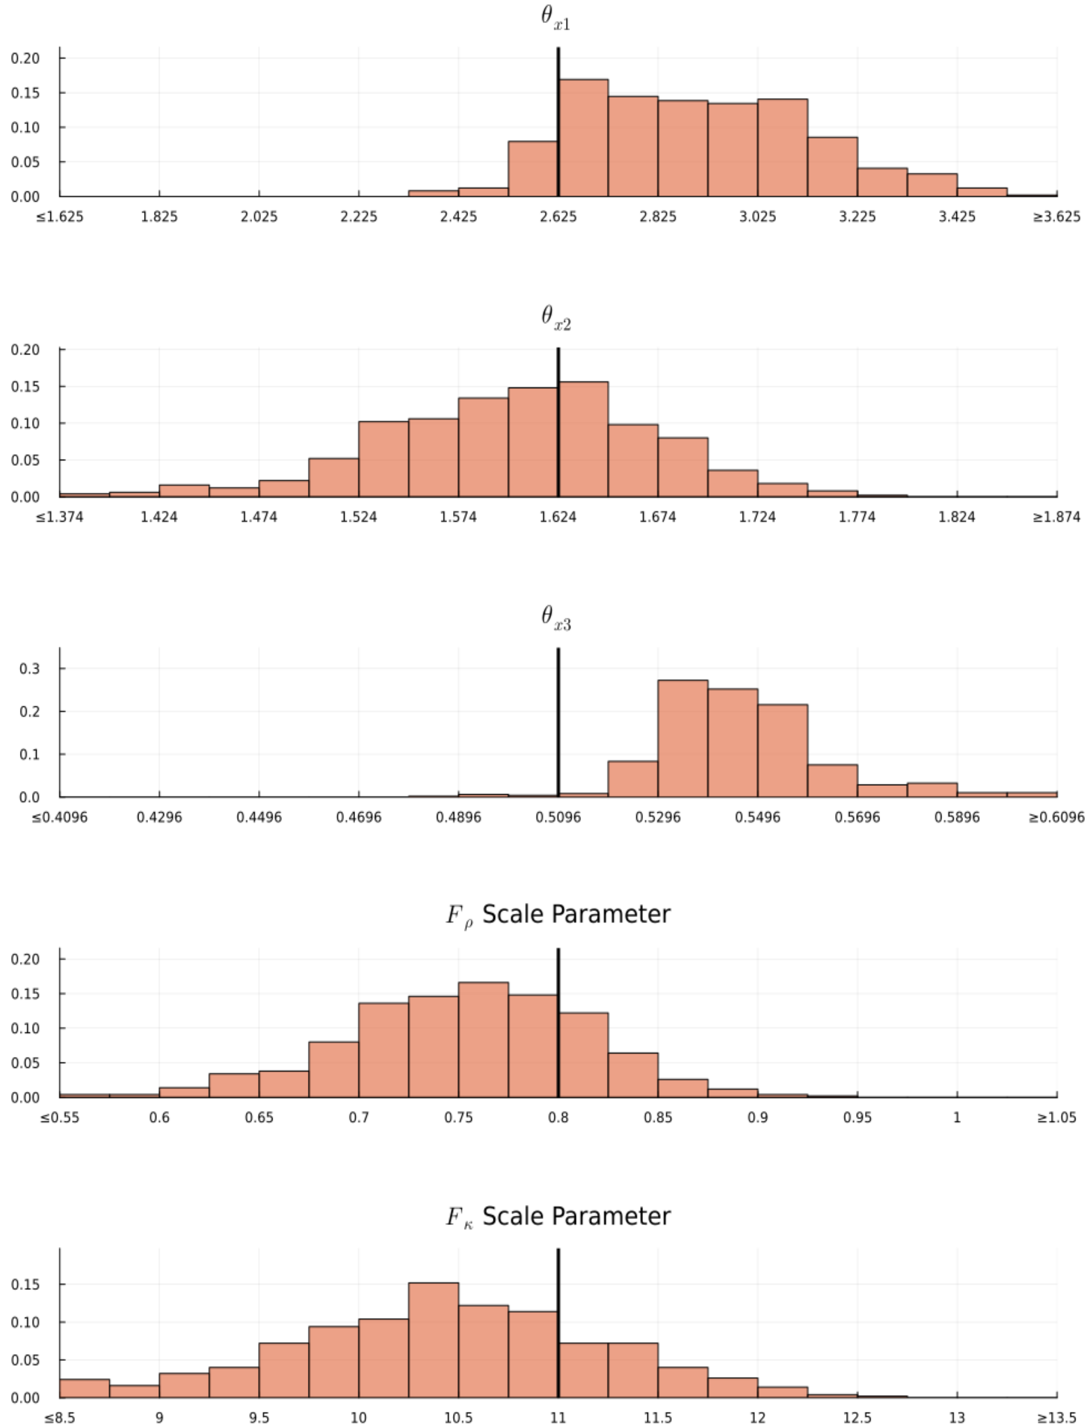

Figure 2: Recursive Estimator – HvB Design

Distribution of the pseudo-MLE recursive estimator across 500 Monte Carlo replications for the HvB design described in Section 4.2. The true parameter values are indicated by the vertical dashed lines.

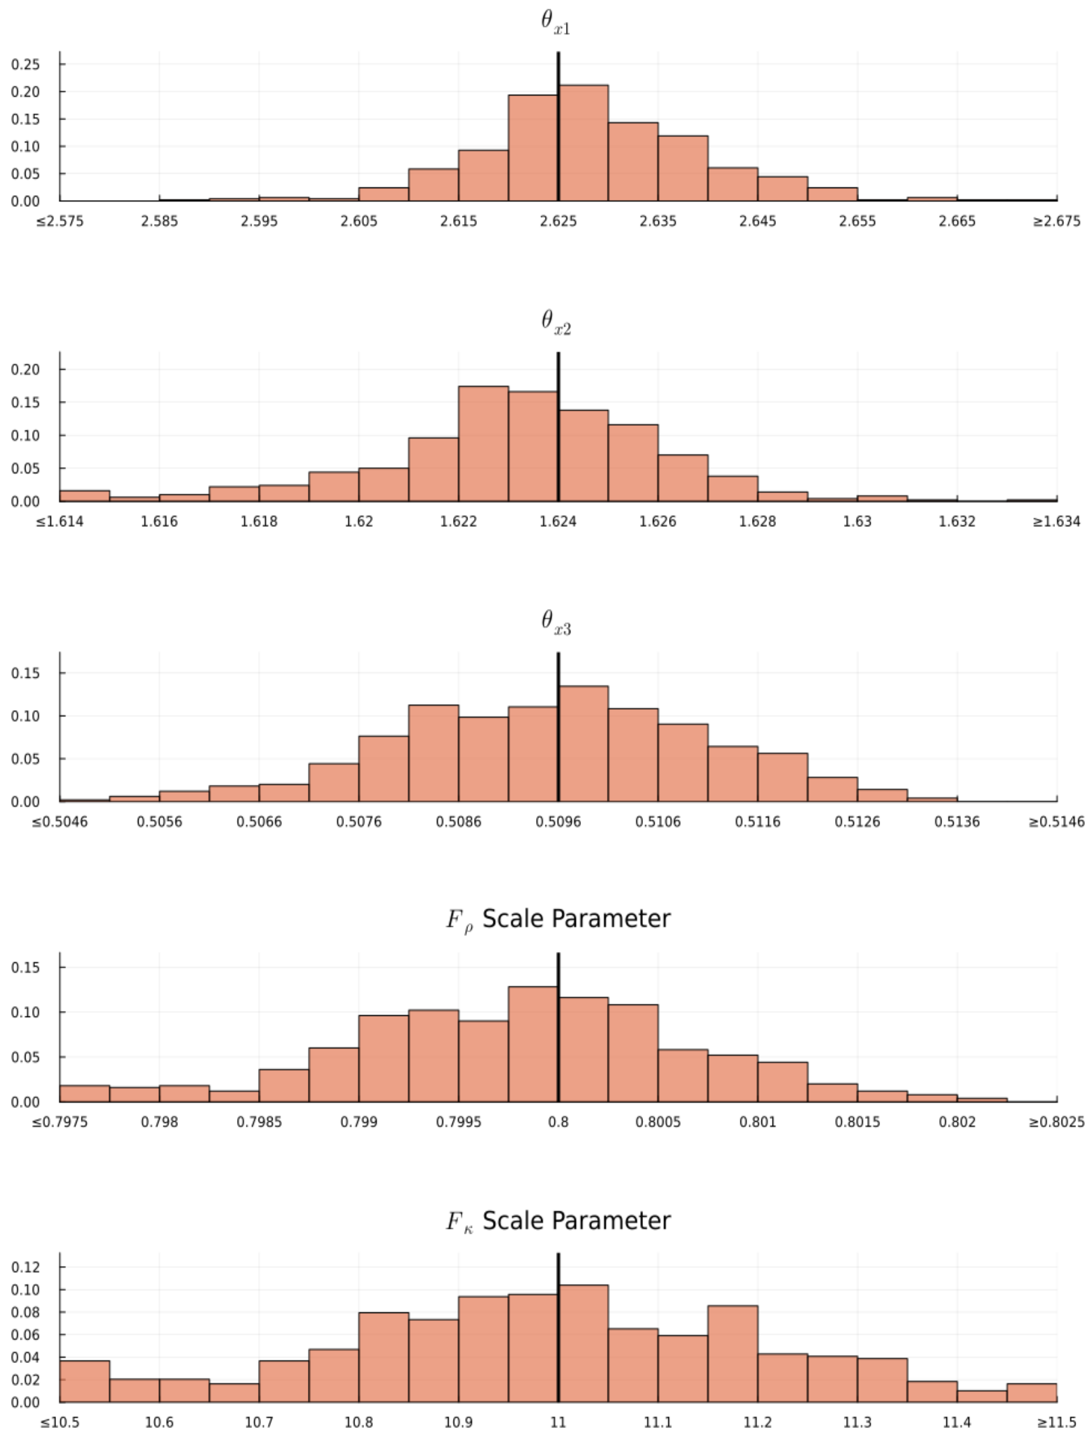

Figure 3: Oracle First Stage – HvB Design

Distribution of the pseudo-MLE recursive estimator across 500 Monte Carlo replications using true MPE policies in the first stage. The true parameter values are indicated by the vertical dashed lines.