

NBER WORKING PAPER SERIES

A COMPUTABLE DYNAMIC OLIGOPOLY MODEL OF CAPACITY INVESTMENT

Gautam Gowrisankaran
Philipp Schmidt-Dengler

Working Paper 32399
<http://www.nber.org/papers/w32399>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2024

We thank El Hadi Caoui, Allan Collard-Wexler, Mitsuru Igami, and seminar participants at the University of Arizona, CREST-INSEE, Universita Cattolica del Sacro Cuore, and Vienna for helpful comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Gautam Gowrisankaran and Philipp Schmidt-Dengler. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

A Computable Dynamic Oligopoly Model of Capacity Investment
Gautam Gowrisankaran and Philipp Schmidt-Dengler
NBER Working Paper No. 32399
May 2024
JEL No. C63,C73,L11,L13

ABSTRACT

This paper analyzes dynamic oligopoly models where investment is the principal strategic variable of interest, there are a large number of investment choices, and there are privately observed shocks to the marginal cost of investment. We show that simulation methods to compute these models can result in non-existence of pure strategy equilibrium. We provide a computationally efficient method to calculate optimal investment probabilities and show how to apply our methods to the recent dynamic empirical literature. The method iteratively finds the investment choices chosen with positive probability and cutoff values of the private information shocks across options in this set.

Gautam Gowrisankaran
Department of Economics
Columbia University
420 West 118th St.
Rm. 1026 IAB, MC 3308
New York, NY 10027
and HEC Montreal and CEPR
and also NBER
gautamg2@gmail.com

Philipp Schmidt-Dengler
University of Vienna
Department of Economics
Austria
philipp.schmidt-dengler@univie.ac.at

1 Introduction

Over the last three decades, researchers have developed a computational literature on dynamic oligopoly models with endogenously heterogeneous firms. These models have allowed researchers to study the evolution of industries in a variety of institutional settings. This has, in turn, yielded important insights into important industrial organization and policy questions. For example, these models have been used to theoretically analyze the evolution of market structure (Pakes and McGuire, 1994), the long-run impact of horizontal mergers (Gowrisankaran, 1999), and the role of learning-by-doing and forgetting and advertising in industry dynamics (Besanko et al., 2010; Doraszelski and Markovich, 2007). This literature has also been used as a framework for empirical work in which the dynamic parameters of a specific model are estimated and policy counterfactuals are then carried out. These approaches have been used to study the dynamics of the hospital (Gowrisankaran and Town, 1997), aircraft (Benkard, 2004), cement (Ryan, 2012), and concrete (Collard-Wexler, 2013) industries, among others.

The canonical model that underlies this literature was developed by Pakes and McGuire (1994) and Ericson and Pakes (1995). In Pakes and McGuire, incumbent firms make quality investments and exit decisions, while potential entrants choose whether to enter. The investment model specifies that each firm invests in order to raise the quality of its product, where the return to the investment is stochastic. The quality of a product can only go up one level in each period. Because of the model's complexity, researchers have typically solved the Markov perfect equilibrium for particular parameters via computation rather than by characterizing the solution analytically.

This paper provides a computable framework to study capacity dynamics in oligopoly models that can also be used in empirical work. In particular, we develop a model with privately observed shocks to the cost of investment. Our model departs from much of the existing literature in the way that cost shocks enter payoffs. In our model, shocks affect payoffs multiplicatively in a function that is increasing in the amount of investment that the firm chooses. Thus, our private information shocks affect the marginal costs of investment rather

than the fixed costs. Our privately observed shocks can also be interpreted as idiosyncratic components of the opportunity cost of capital.

Privately observed payoff shocks have played a key role in the literature in dynamic oligopoly models in particular, and in estimable dynamic discrete choice models more generally. There are two reasons behind this. First, private information shocks are key to establishing existence of pure strategy equilibrium (see Gowrisankaran, 1995; Doraszelski and Satterthwaite, 2010). Second, they generate ex-ante stochasticity in outcomes such as investment. This is useful for models to serve as a basis for empirical work, where one would like to match data with randomness to the predictions of the model (see, for instance, Rust, 1987).

We characterize three types of models that differ from ours in the way they incorporate randomness into the payoffs. First, Pakes and McGuire (1994) and closely related models, allow for *random investment outcomes*. In these models, firms make a continuous investment choice. More investment increases the probability of an upward move in the firm's endogenous state variable (e.g. capacity or quality) by one unit. Firms do not actively retire capacity but instead an exogenous random process governs depreciation. This class of models has been successfully implemented, for instance, to understand horizontal mergers and antitrust policy (Gowrisankaran, 1999; Mermelstein et al., 2020) and to explain empirical facts such as the substantial and persistent firm sizes differences that are observed in most industries (see Besanko and Doraszelski, 2004).

However, this approach to modeling oligopoly capacity dynamics will miss some key features of capacity dynamics. In general, firms often do not change their physical capital, implying that there is large mass point at zero investment. When firms do make changes in capital stock, small negative and positive adjustments are common. However, conditional on making a positive investment, it is also common for those investments to be very large relative to the initial capital stock. For instance, Cooper and Haltiwanger (2006) show that the pattern of firm investment for manufacturing firms has a considerable mass at zero and often includes large and heterogeneous changes in capacity levels.

The second class of models involves *investment fixed cost shocks* with shocks to the fixed

costs of investment and/or disinvestment (Ryan, 2012). This class of models allows firms to choose their optimal level of capacity, possibly by retiring capital. In this class, conditional on a choice of positive or negative investment, the *level* of investment is deterministic. Firms may choose arbitrarily small amounts of investment in order to benefit from favorable investment draws. A deterministic level of investment may also not fit observed data very well.

A third class of models allows for *multinomial cost shocks*. In this class, firms choose from a discrete set of capacity levels, and researchers model the underlying investment decision akin to a probit or multinomial logit setup. They assume that the capacity-level specific shocks are independent across capacity choices. For instance, in their empirical study of the nuclear power plant industry, Rust and Rothwell (1995) model the capacity decision of nuclear power plants with an *i.i.d.* logit shock to each choice of capacity utilization. This approach requires that the number of shocks increases with a finer grid, implying that the grid approximation will have a big impact on the outcomes. It also has the undesirable property that the shock to a large positive investment is independent of a shock to a medium-large positive investment. This type of model also requires the evaluation of multidimensional integrals which tend to be computationally challenging except for the case *i.i.d.* logit errors. Collard-Wexler (2013) models concrete plants in a similar fashion, but considers only three size classes of plant capacity (small, medium, and large). Considering only a small set of choices mitigates the above problem, but at the cost of less good approximations of the true choice set.

Given the issues with these three types of random shocks, researchers more recently have adopted a modeling approach with an additive shock to the marginal cost of investment, in which the cost of investment is monotonic in the level of the shocks. Kalouptside (2018) models a shipyard's decision of how many ships to produce. In addition to a deterministic cost that is linear-quadratic in the number of ships, she models a private cost shock that is proportional to the number of units added. Caoui (2023) studies the transition to digital movie distribution and exhibitions. He models a movie theatre's per-screen adoption cost as the sum of a common price across theaters and a privately observed theater-specific shock. Gowrisankaran et al. (2024) model a utility's investment in generation capacity for different

fuel types. Their cost of capacity investment involves a fixed cost, linear and quadratic terms, and private idiosyncratic shocks to the marginal costs of adding or retiring capacity.

This paper studies computation of the models used in these three applications. In particular, we consider computational dynamic oligopoly models where the principal strategic variable of interest is a discrete investment decision. The model allows for a large and discrete number of potential investment levels that can also approximate a continuous distribution. Our investment process specifies that firms can adjust their capacity as much as desired in a given period by paying the required cost of investment. We model private information shocks to the marginal cost of investment. The shocks are additive to the marginal cost of investment so that the cost of investment is monotone in the level of the shocks. The shocks generate ex-ante stochasticity in the investment outcomes, and ensure existence of equilibrium. Given this model, we develop a simple and efficient computational algorithm to compute exact best responses and equilibria without the need for numerical or Monte Carlo integration methods.

We believe that our paper has two main contributions. First, we characterize the optimal investment policy in models with private information shocks to the marginal cost of investment. We do so without imposing restrictions on the deterministic payoff and transition function. This builds on Kalouptzidi (2018), which relies on convex choice-specific value functions and restrictions on the state transition functions to ensure that the value function is concave in the investment choice, and on Caoui (2023), which imposes a “decreasing difference” restriction on the choice-specific value functions. Allowing for the flexibility of the functional form is important because these assumptions may not necessarily hold in the context of dynamic oligopoly models, where choice-specific value functions may not follow concavity properties.

Second, we show how to compute equilibria of these dynamic oligopoly models without relying on simulation draws. The monotonicity of payoffs in the private information cost shocks implies that investment levels are monotone in the shock, conditional on the state. This monotonicity property then allows us to solve for the cutoff values of these shocks that result across different optimizing levels of investment. Using monotonicity, we further develop

a simple algorithm to identify which investment levels are chosen with positive probability and the optimal cutoffs between different investment levels for any state and firm. These cutoffs fully characterize the optimal choice probabilities given beliefs. The probability of choosing any action is then a continuous function of the future value from each action. Given continuity, existence of equilibrium follows from application of a fixed point theorem. These functions can then be used to solve for a Markov perfect equilibrium of the capacity game in pure investment strategies.

In contrast, we show that an approach where we take a finite number of simulation draws of the private information shock would result in a model that may not have a pure strategy equilibrium. We illustrate this point with simple counterexamples where equilibria do not exist and where increasing the number of simulation draws to any finite level does not solve this problem.

We believe that our general investment framework may be helpful in examining industries where capacity or capital is the main source of heterogeneity. Investment models with quadratic adjustment costs or capital specificity which allow for asymmetric investment behavior have been well-developed in macroeconomics (e.g. Cooper and Haltiwanger, 2006). However, a central hurdle in using this type of model for dynamic *oligopoly* models is computational complexity.¹ A common way of computing equilibria of these models would be via simulation, as is typically done in the macroeconomics literature (see, for example, Cooper and Haltiwanger, 2006; Khan and Thomas, 2008). The approach we develop in this paper allows for the computation of these models, whereas a method that simulated the private information shock may not yield an equilibrium.

The remainder of this paper is divided as follows. Section 2 describes our model. Section 3 discusses computation of equilibria. Section 4 discusses applications from the literature. Section 5 concludes.

¹Our model of investment allows for both quadratic adjustment costs and capital specificity.

2 A Dynamic Framework of Capacity Choice

2.1 Model

We consider capacity investment in a dynamic oligopoly framework with discrete time, $t = 1, 2, \dots, \infty$. Firms discount future payoffs with a discount factor $\beta < 1$. The industry consists of up to N firms at any time t . We denote the set of firms by $\mathcal{N} = \{1, \dots, N\}$, with a typical firm being $i \in \mathcal{N}$. Firms choose from a finite set of capacity levels. Our model is quite general in that capacity in our framework can represent any factor that affects demand. We now describe the states, actions, payoffs, and equilibrium concept, in turn.

Publicly observed states. At time t , there is an L -dimensional vector of publicly observed state variables $s^t \in S \subseteq \mathfrak{R}^L$. The state includes the capacity level of each firm and potentially other firm-specific characteristics. It also includes common characteristics such as aggregate demand or aggregate productivity. We let s_i^t denote the portion of s^t that is affected by the actions of firm i . Similarly, we let s_0^t denote the portion that is unaffected by the actions of any firm, i.e., s_0 evolves exogenously.

Privately observed cost shocks. At time t , each firm i privately observes a real-valued cost shock $\varepsilon_i^t \in \mathfrak{R}$. The shock is not observed by other firms until the end of period t . The shocks ε_i^t are *i.i.d.* and drawn from the strictly monotone and continuous distribution function F . Independence of ε_i^t from the state variables is important, as it allows to integrate over these shocks conditional on the current state.² The strict monotonicity assumption is equivalent to full support on \mathfrak{R} . Continuity of F allows us to ensure existence of pure strategy equilibria.³ The strict monotonicity and continuity assumptions together ensure that multiple investment outcomes are chosen with positive probability *ex ante*. Finally, we assume that $E[\varepsilon_i^t | \varepsilon_i^t \geq \varepsilon] < \infty$ for all ε , which ensures that the expected value conditional on any action is finite.

Actions. After observing the publicly observable state s_t and its own private cost shock ε_i^t , each firm simultaneously chooses next period's capacity. We denote a firm's action by

²Rust (1987) provides a discussion of the independence assumption in dynamic discrete choice models.

³See the discussion in Gowrisankaran (1995) or Doraszelski and Satterthwaite (2010)

a_i^t . We assume that a_i^t is chosen from the action set $A = (\alpha^1, \dots, \alpha^K)$; i.e. A consists of K unique real numbers. We use $o(\cdot)$ to denote the cardinality of sets, so $o(A) = K$. For ease of later analysis and without loss of generality, we impose an increasing order on A : let $\alpha^1 < \alpha^2 < \dots < \alpha^K$. Similarly, we impose an increasing order on the elements of any subset of A . We let an action profile a^t denote the vector of joint actions in period t , $a^t = (a_1^t, \dots, a_N^t) \in A = \times_{i=1}^N A_i$. The cardinality of the action space A is given by K^N .

State transitions. We describe the state transition matrix with a probability density function $g : A \times S \times S \rightarrow [0, 1]$ where a typical element $g(s^{t+1}|a^t, s^t)$ is the probability that state s^{t+1} is reached when the current action profile and state are given by (a^t, s^t) . We require $\sum_{s' \in S} g(s'|a, s) = 1$ for all $(a, s) \in A \times S$. Our framework thus encompasses stochastic depreciation and random investment outcomes as in Pakes and McGuire (1994), as well as a stochastic process for aggregate demand. In addition to encompassing the Pakes and McGuire (1994) framework, our framework allows for firm i 's observable state variable s_i^t to evolve deterministically. In this case, since each firm chooses its next period's capacity level, we can write $s_i^{t+1} = a_i^t$.

Period payoffs. Firm i receives its period payoff at the end of the period, after all actions are observed. We assume that we can additively separate the period payoffs into a deterministic term that does not depend on the private cost shock and a term that becomes stochastic because it is multiplicative in the cost shock. Specifically, we define period payoffs for firm i as a real-valued function defined on $S \times A_i$ and given by:

$$\pi_i(a_i^t, s^t) - c_i(a_i^t, s^t) \times \varepsilon_i^t, \quad (1)$$

where $c_i(\cdot)$ is strictly increasing in its first argument a_i^t , i.e., the firm's stochastic cost of capacity is increasing in the level of capacity chosen. The deterministic term $\pi_i(\cdot)$ depends on the capacity choice for next period a_i^t and the current state of the industry s^t . Conceptually, it includes the profits from selling the product and the deterministic part of investment costs. We assume that both π_i and c_i are bounded: $|\pi_i(\cdot)|, |c_i(\cdot)| < \infty$ for all i .

The discounted sum of future payoffs. Similarly to the period payoffs, the discounted sum of future payoffs consists of the deterministic components and the random cost

component. For firm i , the discounted sum is given by:

$$\mathbb{E} \sum_{\tau=t}^{\infty} \beta^{\tau} [\pi_i(a_i^{\tau}, s^{\tau}) - c_i(a_i^{\tau}, s^{\tau}) \times \varepsilon_i^{\tau}]. \quad (2)$$

The expectation \mathbb{E} is over the realization of ε_i^{τ} as well as own and rival firms' future states and state-contingent actions.

2.2 Markov perfect equilibrium

To analyze equilibrium behavior, we follow Maskin and Tirole (1988) and consider pure *Markovian Strategies* $a_i(s^t, \varepsilon_i^t)$. A strategy for firm i is a function of the firm specific investment cost shock and the publicly observable state variables. The assumption that the profitability shock is independently distributed allows us to write the probability of observing action profile a^t as $Pr(a^t|s^t) = Pr(a_1^t|s^t) \cdots Pr(a_N^t|s^t)$. The Markovian assumption allows us to abstract from calendar time and subsequently we omit the time superscript. It also allows us to define beliefs about each firm's probability of choosing an action at each publicly observable state— $\omega(a, s)$.

Value function. We can define a Bellman equation for firm i given any set of beliefs for the probabilities of actions, ω . Focusing on the ex ante value function, i.e. before the private shock ε_i is realized, we obtain:

$$\begin{aligned} V_i(s|\omega) &= \sum_{k=1}^K \sum_{a \in A \text{ s.t. } a_i = \alpha^k} \omega(a, s) \left\{ [\pi_i(\alpha^k, s) - c_i(\alpha^k, s) \times E(\varepsilon_i | a_i = \alpha^k, s)] \right. \\ &\quad \left. + \beta \sum_{s' \in S} V_i(s'|\omega) g(s'|a, s) \right\}. \end{aligned} \quad (3)$$

Here, E denotes the expectation operator with respect to the firm's investment cost shock. The finiteness of the action and state space guarantees the existence of the value function $V_i(s|\omega_i)$ in equation (3).

Choice-specific values. To further expost the optimal choices of investment necessary to characterize and compute the equilibrium of our model, we follow Hotz and Miller (1993) and define the *choice-specific value* :

$$v_i^k(s|\omega_i) = \pi_i(\alpha^k, s) + \beta \sum_{a_{-i} \in A_{-i}} \sum_{s' \in S} g(s'|s, k, a_{-i}) V_i(s'|\omega_i), \quad (4)$$

as firm i 's value net of the random component of cost $c_i(a_i, s) \times \varepsilon_i$ when it chooses action $a(s, i) = k$. It is then optimal to choose action $a(s, i) = k$ under beliefs ω_i whenever

$$v_i^k(s|\omega_i) - c_i(\alpha^k, s) \times \varepsilon_i \geq v_i^\ell(s|\omega_i) - c_i(\alpha^\ell, s) \times \varepsilon_i, \forall \ell \neq k. \quad (5)$$

This characterizes the optimal decision rule up to a set of measure zero. For this zero measure set we assume, without loss of generality, that whenever equation (5) holds with equality, the firm chooses the higher action. The optimal policy $a_i(\varepsilon_i, s)$ then satisfies:

$$a_i(\varepsilon_i, s) = \operatorname{argmax}_{\alpha^k \in A} \{v_i^k(s|\omega_i) - c_i(\alpha^k, s) \times (\varepsilon_i)\}. \quad (6)$$

The probability that firm i chooses action $a_i = \alpha^k$ in state s thus given by

$$\begin{aligned} p_i^k(s) &= \psi_i^k(s|\omega_i) \\ &= \Pr \left(\begin{array}{l} v_i^k(s|\omega_i) - c_i(\alpha^k, s) \times \varepsilon_i \geq \\ v_i^\ell(s|\omega_i) - c_i(\alpha^\ell, s) \times \varepsilon_i, \forall \ell \neq k \end{array} \right). \end{aligned}$$

This relationship holds for all firms $i \in \mathcal{N}$ and states s , and every action a . Without loss of generality, we set the lowest capacity choice $a_i = \alpha^1$ to be the reference action whose probability $p_i^0(s)$ is given by $1 - \sum_{k=2}^K p_i^k(s)$. This results in a system of $L \cdot N \times (K - 1)$ equations, which we can write compactly in vector notation as

$$p = \psi(\omega). \quad (7)$$

where p denotes the $L \cdot N \times (K - 1)$ -dimensional vector of choice probabilities and ω the $L \cdot N \times (K - 1)$ -dimensional vector of beliefs.

A Markov perfect equilibrium (MPE) is a set of strategies and beliefs regarding the probabilities of actions for each publicly observable state,

$$(a, \omega) = (a_1, \dots, a_N; \omega_1, \dots, \omega_N),$$

that satisfies the following conditions. First, each firm's strategy $a_i(s, \varepsilon_i)$ is Markovian and a best response to a_{-i} given beliefs ω_i . Second, beliefs about each firm's probability of choosing an action at each publicly observable state— $\omega_i(a, s)$ —are consistent with strategies a .

In a MPE, it must hold that beliefs ω have to correspond to choice probabilities p so that equation (7) becomes

$$p = \psi(p) \tag{8}$$

It follows that any p satisfying (8) constitutes an equilibrium. Note that $\psi(\cdot)$ is a mapping from a $L \cdot N \times (K - 1)$ -dimensional unit simplex into itself. Since $\psi(\cdot)$ is continuous in p , Brouwer’s fixed point theorem applies and (7) it has a solution in p . Note further that when payoffs are symmetric, a symmetric MPE exists. This is because we can use the same argument, but restricting to symmetric strategies in the presence of symmetric payoffs.

3 Computation of equilibrium

3.1 Simulation-based computation and nonexistence of equilibria

In order to compute a dynamic equilibrium of the model, one can use the method of successive approximations. This method is closely adapted from Pakes and McGuire (1994) and related studies. The idea is to repeatedly compute the mapping ψ until a fixed point is reached.

Specifically, one can start with beliefs of rivals’ choice probabilities, ω_i and the corresponding ex-ante value function, $V_i(s|\omega_i)$, and beliefs of rivals’ choice probabilities, ω_i . The algorithm updates these matrices as follows. For each state (s, i) at each iteration, it first computes a choice-specific value function, $v_i^k(s|\omega_i)$, for $k = 1, \dots, K$, by applying (4) using the previous iteration of the ex-ante value function and beliefs. Using the choice-specific values, it then solves for firm i ’s optimal policies conditional on ε_i , $a_i(\varepsilon_i, s)$, using (6). The algorithm then integrates over ε_i to solve for the ex-ante value function $V_i(s|\omega_i)$, as in (3) and uses the new $V_i(s|\omega_i)$ to update transition probability beliefs ω_i . It performs this calculation for all states, iterating on these two steps until a convergence criterion—based on distance between subsequent iterations being close to 0—is satisfied.

The one central complication with this algorithm is the necessity of integrating over the ε draws to calculate the ex-ante value function. A standard approach would be to calculate this integral by simulation. With the standard approach, one would simulate over a finite

number of draws for ε . The problem with the standard approach is that a pure strategy equilibrium for the model with a finite number of draws may not exist even when one exists for the limiting model.

To understand the lack of existence, recall that existence of equilibrium in our model relies on the continuity of the mapping $\psi(\cdot)$. Yet, for the model approximated via simulation, the probability of any action will be discontinuous in valuations because it is the sum of the probabilities over a finite number of draws, each of which has one associated optimal action. We show this by two examples that add private information to textbook games.

Matching pennies. Consider the following (static) matching pennies game, a classic example of a game where no pure strategy Nash equilibrium exists. We can formulate matching pennies as a special case of our model, where there is only one state, so capacity $s_i = 1$ for both players, and players simultaneously choose whether to increase capacity to 2. The deterministic component of firm 1's payoff for action $a_1 = \alpha^1$ is -0.75 when both choose the same action, and 1.25 otherwise. The deterministic component of firm 2's payoff is 1 when both choose the same action, and -1 otherwise. Our model solves the non-existence by adding a private information shock to the payoff when choosing action $a_i = \alpha^2$.⁴ For simplicity, we consider uniformly distributed shocks with $\varepsilon_i \sim U(-0.5, 0.5)$. The modified game is given below.

		Player 2	
		$a_2 = \alpha^1$	$a_2 = \alpha^2$
Player 1	$a_1 = \alpha^1$	$(-0.75, 1)$	$(1.25, -1 - \varepsilon_2)$
	$a_1 = \alpha^2$	$(1 - \varepsilon_1, -1)$	$(-1 - \varepsilon_1, 1 - \varepsilon_2)$

Let p_i be Player i 's probability of choosing action $a_i = 1$. Player 1 chooses action $a_1 = 1$ if and only if $\varepsilon_1 > 4p_2 - 2.25$. Player 2 chooses action $a_1 = 1$ if and only if $\varepsilon_2 > 2 - 4p_2$. The modified game has a unique pure strategy equilibrium. Player 1 plays action 2, if and only if $\varepsilon_1 > -1/68$. Similarly, Player 2 plays action 2 if and only if $\varepsilon_2 > -1/17$. First suppose

⁴This matching pennies model corresponds to a capacity adjustment model where there is a stochastic cost when the firm needs to adjust capacity to 2.

one tried to solve for an equilibrium by simulation. Consider an approximation model with two draws for each player with realizations $\{-0.3, 0.3\}$. A pure strategy equilibrium exists, because each player would play each action with probability one half, choosing action 1, when drawing $\varepsilon_i = -0.3$ and action 2 otherwise.

Now suppose that the draws are $\{-0.2, 0.2\}$. We considered the cases where one or both players choose the same action across the two draws, i.e., $p_i \in \{0, 1\}$, for some i . We verified that none of these cases forms an equilibrium. It must then be that in any equilibrium, $p_i \in (0, 1)$ for $i = 1, 2$. By the nature of the private information shock, it further follows that, in a pure strategy equilibrium, Player 2 again chooses action 1, when drawing $\varepsilon_i = -0.2$ and action 2 otherwise. This implies that $p_2 = 0.5$. Because Player 1 chooses action $a_1 = 1$ if and only if her draw was larger than $4p_2 - 2.25 = 0.25$, she would never choose action 1, which implies that she chooses the same action $a_1 = 2$ for both realizations, which we know cannot be an equilibrium. Therefore, no pure strategy equilibrium exists in this case. A successive approximation algorithm would cycle forever. This shows that a pure strategy equilibrium does not exist for some values of the draws for the approximated model with discretized shocks.

Duopoly exit. Consider a simple game where firms only consider whether to stay in the market or exit. This corresponds to variant of our capacity game with a discount factor of 0, possible actions being a capacity of either $\alpha^1 = 0$ or $\alpha^2 = 1$, and period payoffs that are affected immediately by the exit choices. If one firm remains active, the firm earns a deterministic component of profits equal to $\bar{\pi}(1)$, if two firms remain active, they each earn duopoly profits $\bar{\pi}(2)$. Firms that have exited earn profits of zero. Assume that $\bar{\pi}(2) < 0 < \bar{\pi}(1)$, so that, with a private information shock, there is no pure strategy symmetric equilibrium, but two asymmetric pure strategy equilibria (0,1) and (1,0). We now add an *i.i.d.* cost shock ε drawn from distribution function F to be paid if a firm remains active (i.e. $c(\alpha^k) = \alpha^k$, for $k \in \{1, 2\}$). Existence of pure strategy equilibrium can easily be established for this model. A symmetric pure strategy equilibrium corresponds to the choice probability

g of a firm remaining active, defined by the solution to the equation

$$g = F(g\bar{\pi}(2) + (1 - g)\bar{\pi}(1)).$$

Since, in our model, F is strictly increasing with positive support over the entire real line, at least one such equilibrium is guaranteed to exist. Now a simulation approach could not use the above equation, but instead would draw shocks from the discretely approximated distribution F . Suppose that one chooses two values ε_1 and ε_2 , each drawn with probability $1/2$. Without loss of generality, assume that $\varepsilon_1 < \varepsilon_2$. A pure strategy for each firm would be to always exit (implying choice $g = 0$), remain active whenever the cost shock is low ($g = 1/2$), or always remain active ($g = 1$).

Symmetric equilibria only exist for some sets of values of the private information draws. For instance no symmetric equilibrium exists for the four following sets of values:

$$\begin{aligned} & \{\varepsilon_1 < \bar{\pi}(2) < \varepsilon_2 < (1/2)(\bar{\pi}(1) + \bar{\pi}(2))\}, \\ & \{\bar{\pi}(2) < \varepsilon_1, \varepsilon_2 < (1/2)(\bar{\pi}(1) + \bar{\pi}(2))\}, \\ & \{(1/2)(\bar{\pi}(1) + \bar{\pi}(2)) < \varepsilon_1, \varepsilon_2 < \bar{\pi}(1)\}, \\ & \{(1/2)(\bar{\pi}(1) + \bar{\pi}(2)) < \varepsilon_1 < \bar{\pi}(1) < \varepsilon_2\} \end{aligned}$$

Consider for instance the first set of values. Whenever the other firm remains active with probability 0 or $1/2$ it is optimal to always remain active, i.e. $g = 1$. However, the best response to the rival remaining active with probability 1 is to remain active whenever the cost shock realization is low. This shows that a pure strategy symmetric equilibrium does not exist for some values of the draws for the approximated model with discretized shocks.

3.2 Efficient computation

Our innovation is the development of an algorithm that allows us to find the exact probabilities of choosing each capacity level α^k conditional on choice-specific values for a state (s, i) . Our algorithm does not use simulation and is quick to compute, thereby making it suitable to use nested within the dynamic game solution.

This subsection considers the computation of the probability of form i and each action given the state s . To ease notation, in this subsection, we use $a(\varepsilon)$, $c(\alpha^k)$, and v^k to refer to $a_i(\varepsilon|s)$, $c_i(\alpha^k, s)$, and $v_i^k(s|\omega_i)$ respectively. We start with the following lemma:

Lemma 1. *The action function $a(\varepsilon)$ is weakly decreasing in ε .*

Proof This proof follows from Topkis Theorem⁵ but we prove it directly. Consider two actions k and j , with $k > j$;

$$\begin{aligned}
 & j \text{ strictly preferred to } k \\
 \iff & v^j - c(\alpha^j) \times \varepsilon > v^k - c(\alpha^k) \times \varepsilon \\
 \iff & \varepsilon > \frac{v^k - v^j}{c(\alpha^k) - c(\alpha^j)},
 \end{aligned} \tag{9}$$

where the third line uses the fact that $c(\cdot)$ is increasing, implying that $c(\alpha^k) - c(\alpha^j)$ is positive. From (9), for any two actions, the higher action will only be chosen with lower ε , implying that a is weakly decreasing in ε . ■

Using Lemma 1, we define the ε cutoff between any two choices.

Definition For $1 \leq j < k \leq K$, let the “ ε cutoff” be

$$\bar{\varepsilon}(j, k) = \frac{v^k - v^j}{c(\alpha^k) - c(\alpha^j)}.$$

This definition and Lemma 1 lead directly to another (small but important) result:

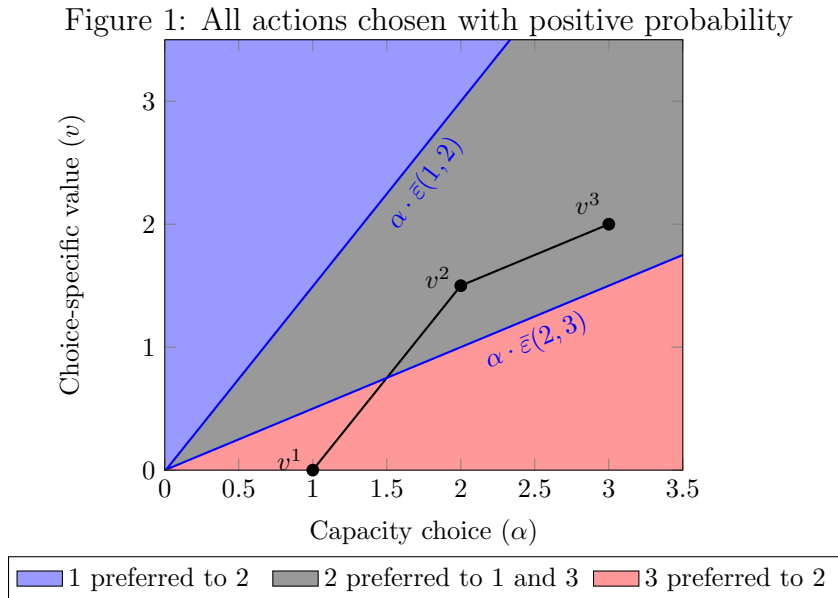
Lemma 2. *For $1 \leq j < k \leq K$, the firm strictly prefers action j to $k \iff \varepsilon > \bar{\varepsilon}(j, k)$.*

Proof This result follows directly from the proof of Lemma 1.

Sometimes—but not always—one can directly use the ε cutoffs to define the probabilities of each action. Before turning to the general results, we illustrate this with an example. We consider the case of three possible actions, $K = 3$, and linear capacity costs, $c(\alpha^k) = k$. The

⁵See for instance Theorem 1 in Amir (2005).

payoff from an action is therefore $v^k - a \times \varepsilon$. Let the smallest action $k = 1$ yield a choice-specific value of $v^1 = 0$ and the largest action $k = 3$ yield a choice-specific value of $v^3 = 2$. In Figure 1, we consider the case where the middle action $k = 2$ comes with a choice-specific value of $v^2 = 1.5$. Here, “ ε cutoffs” are $\bar{\varepsilon}(1, 2) = 1.5$ and $\bar{\varepsilon}(2, 3) = 0.5$, respectively. The “ ε cutoffs” determine the slopes of the dark blue lines, which are by construction parallel to the corresponding black lines that connect the choice-specific values. Lemma 2 implies that $a(\varepsilon) = 1$ whenever $\varepsilon > 1.5$. Further, $a(\varepsilon) = 2$ when $0.5 < \varepsilon < 1.5$, and $a(\varepsilon) = 3$ when $\varepsilon < 0.5$. Consequently, all three actions are chosen with positive probability.⁶ Action $a_i = 1$ with probability $1 - F(1.5)$, $a_i = 2$ with probability $F(1.5) - F(.5)$, and $a_i = 3$ with probability $F(.5)$. We will now see that the key is that the cutoffs are declining in the action: $\bar{\varepsilon}(1, 2) = 1.5 > \bar{\varepsilon}(2, 3) = 0.5$, which corresponds to the conditional choice value function being discrete concave.

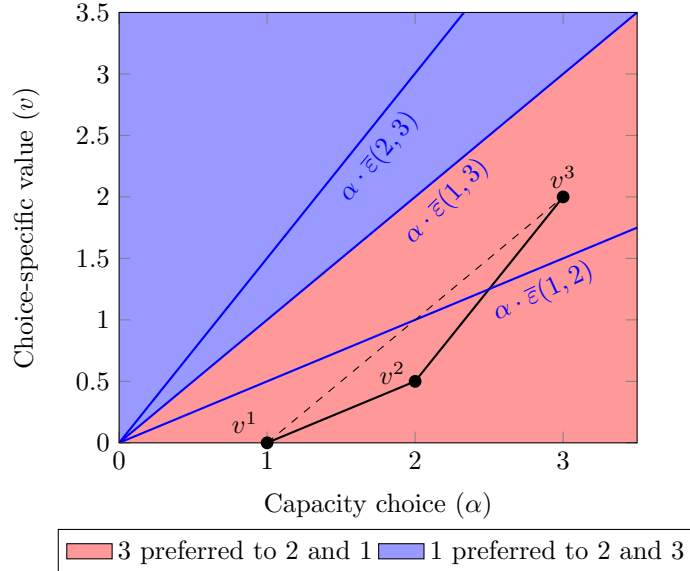


In Figure 2, we alter the choice-specific value of the second action $k = 2$ to $v^2 = .5$. In this case, v^2 lies below any convex combination of v^1 and v^3 , which we show with the dashed line. Not coincidentally, the cutoffs have flipped order, with $\bar{\varepsilon}(1, 2) = 0.5$ and $\bar{\varepsilon}(2, 3) = 1.5$. Also not coincidentally, using Lemma 2, one can verify that $a(\varepsilon) = 1$ when $\varepsilon > 2$, $a(\varepsilon) = 3$

⁶Note that in Section 2.2 we adopted the convention that at the cutoff the firm chooses the larger action.

when $\varepsilon < 2$. Consequently, $a(\varepsilon) = 2$ is never optimal.

Figure 2: Only two actions chosen with positive probability



Overall then, Lemma 2 and Figures 1 and 2 imply that the presence of decreasing differences in choice-specific values guarantees the monotonicity of equilibrium strategies, but does not guarantee the concavity of the choice-specific value functions. Consequently, it does not imply that each action is played with positive probability in equilibrium.

Focusing specifically on Figure 2, actions are chosen with positive probability when the ε cutoffs with respect to them are never reversed in order, or equivalently, when their choice-specific values lie above any convex combination of choice-specific values for choices above and below them. We formalize:

Proposition 1. *Let $C(A) \subseteq A$ be the set of actions that are chosen with positive probability.*

The following three statements are equivalent:

- (i) *Action k satisfies $k \in C(A)$*
- (ii) *$\nexists j, l \in A$ such that $j < k < l$ and $\bar{\varepsilon}(k, l) \geq \bar{\varepsilon}(j, k)$*
- (iii) *$\forall j, l \in A$ such that $j < k < l$ and $v^k > \lambda v^j + (1 - \lambda)v^l$ where $\lambda = \frac{c(\alpha^k) - c(\alpha^j)}{c(\alpha^l) - c(\alpha^j)}$.*

Proof We consider first the equivalence of the first two statements. We show (the contrapositive of) the first statement implying the second one, that if $\exists j, l \in C(A)$ s.t. $j < k <$

l and $\bar{\varepsilon}(k, l) \geq \bar{\varepsilon}(j, k)$ then $k \notin C(A)$. If such j, l exist, then for any $\varepsilon > \bar{\varepsilon}(j, k)$, the firm will not pick k since it prefers j to k by Lemma 2. Any $\varepsilon < \bar{\varepsilon}(j, k)$ also satisfies $\varepsilon \leq \bar{\varepsilon}(k, l)$, so the firm will also not pick k in this case since it prefers l to k . Thus, the firm picks k , if at all, at $\bar{\varepsilon}(j, k)$ which has probability 0.

Now, we show that the second statement implies the first. Suppose $\forall j, l \in A$ with $j < k < l$, $\bar{\varepsilon}(k, l) < \bar{\varepsilon}(j, k)$. Let $\bar{\varepsilon}^L(k) = \max\{-\infty, \max_{l>k}\{\bar{\varepsilon}(k, l)\}\}$ and let $\bar{\varepsilon}^U(k) = \min\{\infty, \min_{j<k}\{\bar{\varepsilon}(j, k)\}\}$. Now consider $\varepsilon \in [\bar{\varepsilon}^L(k), \bar{\varepsilon}^U(k)]$. For any such ε , k is preferred to j , $\forall j < k$ and k is preferred to l , $\forall l > k$, and thus the firm chooses k . Moreover, this interval is non-empty by assumption. Thus, k is chosen with positive measure or probability.

Finally, we show that the second and third statements are equivalent.

$$\begin{aligned} & \bar{\varepsilon}(k, l) < \bar{\varepsilon}(j, k), \quad \forall j < k < l \\ \iff & \frac{v^l - v^k}{c(\alpha^l) - c(\alpha^k)} > \frac{v^k - v^j}{c(\alpha^k) - c(\alpha^j)}, \quad \forall j < k < l \\ \iff & v^k > \lambda v^j + (1 - \lambda)v^l, \quad \forall j < k < l. \quad \blacksquare \end{aligned}$$

Proposition 1 shows that each action is chosen with positive probability if and only if it is in a discrete analog of the concave hull with respect to other actions, where concavity is defined only over the discrete actions and using the measure λ . The proposition also implicitly provides an algorithm for evaluating which actions are chosen with positive probability. Moreover, this algorithm can be used to compute the region where ε is chosen: evaluating $(\bar{\varepsilon}^L(k), \bar{\varepsilon}^U(k))$, if this forms an interval with positive mass, $k \in C(A)$, and is chosen over this interval.

While this algorithm would correctly compute the probabilities of each action, we show that it is not the most efficient. In particular, it requires computing $\bar{\varepsilon}(j, k)$ for each $1 \leq j < k \leq K$, of which there are $K(K - 1)/2$. Thus, the number of calculations grows with the square of K . However, we develop a computational algorithm that requires only a number of calculations that grows linearly with K . Our algorithm works by iteratively constructing the set of actions chosen with positive probability (those in $C(A)$) and then comparing each element against its neighbor in the set that ultimately becomes $C(A)$. Recall that speed of computation here is crucial since this computation process must be repeated for each industry

state at each stage of the model. We offer the following:

Proposition 2. *Consider the following iterative algorithm:*

1. Initialize $\hat{C} = A$ and $k = 2$.
 - At each step, the algorithm considers a \hat{C} and k .
 - \hat{C} is the candidate $C(A)$, while k and its neighbors are the elements being considered for exclusion from the discrete concave hull.
2. If $\bar{\varepsilon}(\hat{C}_{k-1}, \hat{C}_k) > \bar{\varepsilon}(\hat{C}_k, \hat{C}_{k+1})$, then:
 - If $k = o(\hat{C}) - 1$, exit the algorithm with \hat{C} .
 - Otherwise, go back to the beginning of Step 2 of the algorithm with \hat{C} and $k + 1$.
3. If $\bar{\varepsilon}(\hat{C}_{k-1}, \hat{C}_k) \leq \bar{\varepsilon}(\hat{C}_k, \hat{C}_{k+1})$, then:
 - Drop \hat{C}_k so that $\hat{C} = (\hat{C}_1, \dots, \hat{C}_{k-1}, \hat{C}_{k+1}, \dots, \hat{C}_{o(\hat{C})})$.
 - If $o(\hat{C}) = 2$, exit the algorithm with \hat{C} .
 - Otherwise, go back to (the beginning of) Step 2 of the algorithm, using the following values:
 - If $k > 2$, use \hat{C} and $k - 1$.
 - If $k = 2$, use the new \hat{C} and k .

Then, the output of this algorithm satisfies $C(A) = \hat{C}$.

Proof First, we show (the contrapositive of) the property that any element $k \in \hat{C}$ satisfies $k \in C(A)$. If an element $k \notin C(A)$, then it was removed at some stage of the algorithm because of ε cutoffs that are in the wrong order against two neighbors at the time. But, Proposition 1 shows that each element $k \in \hat{C}$ satisfies the ε cutoffs being in the right order against all other choices. Hence, $k \notin \hat{C}$.

To prove that any element $k \in C(A)$ satisfies $k \in \hat{C}$, we first show the following claim: $\bar{\varepsilon}(C(A)_k, C(A)_{k+1}) < \bar{\varepsilon}(C(A)_{k-1}, C(A)_k)$ for $k = 2, \dots, o(C(A)) - 1$.⁷ Suppose, by contradiction, that $\exists k$ such that $\bar{\varepsilon}(C(A)_k, C(A)_{k+1}) \geq \bar{\varepsilon}(C(A)_{k-1}, C(A)_k)$. Then, by construction, the algorithm compares each k against its neighbors in $C(A)$ at some point in the algorithm. Thus, at this point, $C(A)_k$ would have been dropped from the algorithm.

We next show a second claim: for any $k \notin C(A)$, $\exists j, l \in C(A)$ such that $\bar{\varepsilon}(j, k) \leq \bar{\varepsilon}(k, l)$. Suppose, again by contradiction, that $\exists k \notin C$ such that $\bar{\varepsilon}(j, k) > \bar{\varepsilon}(k, l), \forall j, l \in C(A)$. Then, consider

$$\varepsilon \in \left(\max_{l \in C(A), l > k} \{\bar{\varepsilon}(k, l)\}, \min_{j \in C(A), j < k} \{\bar{\varepsilon}(j, k)\} \right).$$

By the contradictory assumption, this interval has positive mass. Moreover, for any ε in this interval, k is preferred to all $j < k$ if $j \in C(A)$ and to all $l > k$ if $l \in C(A)$. This implies that within this interval, either k or some other $k' \notin C(A)$ are chosen throughout. This then contradicts the fact that elements $k \notin C(A)$ are chosen with zero probability.

Using these two claims, by contradiction, consider now the possibility that for some $C(A)_k$, $C(A)_k \notin \hat{C}$. By the second claim, $\exists j', l' \in C(A)$ such that $\bar{\varepsilon}(C(A)_k, l') \leq \bar{\varepsilon}(j', C(A)_k)$. By the first part of the proposition, $j', l' \in C(A)$. Denote these elements $C(A)_j$ and $C(A)_l$ respectively. Now, by the first claim,

$$\begin{aligned} \bar{\varepsilon}(C(A)_{l-1}, C(A)_l) &< \dots < \bar{\varepsilon}(C(A)_{k+1}, C(A)_k) \\ &< \bar{\varepsilon}(C(A)_{k-1}, C(A)_k) < \dots < \bar{\varepsilon}(C(A)_j, C(A)_{j+1}). \end{aligned}$$

Moreover, note that $\bar{\varepsilon}(C(A)_{k-1}, C(A)_k) \leq \bar{\varepsilon}(C(A)_j, C(A)_k) \leq \bar{\varepsilon}(C(A)_j, C(A)_{j+1})$. If not, suppose $\bar{\varepsilon}(C(A)_j, C(A)_k) < \bar{\varepsilon}(C(A)_{k-1}, C(A)_k)$. Then, for ε in a neighborhood immediately to the right of $\bar{\varepsilon}(C(A)_j, C(A)_k)$, $C(A)_j$ is preferred to $C(A)_k$. But, at this point, $C(A)_k$ is preferred to $C(A)_{k-1}$, which is preferred to $C(A)_{k-2}$ etc. and is ultimately preferred to $C(A)_j$, by the first claim. Thus, this yields a contradiction. Similarly, if $\bar{\varepsilon}(C(A)_j, C(A)_k) > \bar{\varepsilon}(C(A)_j, C(A)_{j+1})$ there would be an equivalent contradiction. By an analogous argument, $\bar{\varepsilon}(C(A)_{l-1}, C(A)_l) \leq \bar{\varepsilon}(C(A)_k, C(A)_l) \leq \bar{\varepsilon}(C(A)_k, C(A)_{k+1})$. Thus, $\bar{\varepsilon}(C(A)_k, C(A)_l) < \bar{\varepsilon}(C(A)_j, C(A)_k)$ which contradicts the initial assumption. Thus, $C(A) = \hat{C}$.

⁷We use $C(A)_k$ to denote the k th element of $C(A)$.

To see that each interior $C(A)_k$ is preferred exactly when $\varepsilon \in [\bar{\varepsilon}(C(A)_k, C(A)_{k+1}), \varepsilon(C(A)_{k-1}, C(A)_k)]$, note that it cannot be preferred outside this region, because the neighboring elements in $C(A)$ are preferred to it. Moreover, no element in $C(A)$ is preferred to it in this region, and no element outside $C(A)$ is ever chosen. An analogous argument holds for $C(A)_1 = a_1$ and $C(A)_{o(C(A))} = a_K$. ■

Proposition 2 defines the ranges of ε for which each action is chosen, while simultaneously creating the set $C(A)$. Combining these insights, we offer:

Corollary 1. *A given action $1 < k < K$ is weakly preferred over all other actions $\iff \varepsilon \in [\bar{\varepsilon}(C(A)_k, C(A)_{k+1}), \bar{\varepsilon}(C(A)_{k-1}, C(A)_k)]$, with action 1 preferred for $\varepsilon \in [\bar{\varepsilon}(C(A)_1, C(A)_2), \infty)$ and action K preferred for $\varepsilon \in (-\infty, \bar{\varepsilon}(C(A)_{o(C(A))-1}, C(A)_{o(C(A))})]$.*

As we discussed above, a principle value of Proposition 2 is that it shows that the computation time for optimal strategies for our algorithm is linear in the cardinality of the action set:

Corollary 2. *Using the algorithm defined by Proposition 2, the number of ε cutoffs that are computed is at least $K - 1$ and at most $2K - 3$.*

Proof At each iteration of the algorithm, either k increases by 1 or an element is removed and the right-most element stays the same. The worst case for computation is that k goes from 2 to K and then all interior $K - 2$ elements are dropped. This worst case requires the computation of $2K - 3$ cutoff values. The best case, which drops no elements, requires the computation of $K - 1$ cutoff values. ■

The limitation in cardinality is key in making our algorithm feasible to compute. Recall that the cardinality is important because the algorithm outlined in Proposition 2 will be performed many times at every state when solving for the MPE with a dynamic oligopoly.

4 Applications From the Literature

Many papers in the dynamic oligopoly literature build on the Pakes and McGuire (1994) (and Ericson and Pakes, 1995) model. However, there are important differences between

this model and ours. We thus first discuss how our results relate to Pakes and McGuire (1994) and why the differences imply that our method may not simplify the computation of the original model. We then turn to three recent dynamic empirical industrial organization papers inspired by Pakes and McGuire (1994), but where firms have private information shocks and can adjust capacity through making one of a large set of potential decisions. We show the mapping between our model and each of these papers, which also demonstrates that our results may be useful in these settings.

4.1 Quality ladder model in Pakes and McGuire (1994)

Pakes and McGuire (1994) model a differentiated product quality ladder game with entry and exit. The publicly observable state vector s^t indicates the quality level of each firm relative to the outside option. Each period, firms invest a continuous amount $a_i \geq 0$ with the goal of increasing product quality s_i^t by 1. The probability that a firm's product quality increases by 1 is increasing in the investment amount. Each period, following the investment phase, firms simultaneously set prices and earn profits. The price-setting decisions are the same as in a static game and hence firms earn Bertrand differentiated-products profits, based on these strategies.

Our model relates to this game, but does not fit exactly into their framework, for a couple of reasons. First, our actions a_i^t are discrete while the action set in Pakes and McGuire is continuous. Focusing on the investment decision, we can approximate their game as a version of our game with a large number of discrete investment choices. The period return becomes

$$\bar{\pi}_i(s_t) - a_i^t c, \tag{10}$$

where the actions span a wide range of positive levels.

Second, in Pakes and McGuire, there is no privately observed cost shock. This is true even in the discretized version in (10). We can see this by noting that (10) would be similar to of our period profits shown in (1) if we added the final term with ε . Pakes and McGuire do not need this term for a pure strategy Markov perfect equilibrium, because they obtain stochasticity (and hence continuity in values) from the random realization of investment,

rather from random investment cost. Therefore, in Pakes and McGuire, quality investments are deterministic conditional on observable state variables while investment outcomes are not. Without the private information shock, applying our method would not reduce computational time. Thus, to the extent that researchers use the exact Pakes and McGuire specification without a private information shock, there is no need to apply our technique. Indeed, computation of Pakes and McGuire involves a different exact solution for the level of investment at each state (Pakes et al., 1992).

Third, Pakes and McGuire (1994) model permanent entry and exit. Our model can be modified to incorporate permanent entry and exit. In this case, we can consider a firm at the lowest capacity level to be one that has exited. An active firm that chooses to exit would expect to receive no future payoffs upon exit, except for a one time scrap value. An exited firm could then be replaced by a future entrant that starts at the lowest capacity level.⁸

Many empirical papers have built on Pakes and McGuire (1994) to model capacity and other similar attributes. In these contexts, firms make infrequent but large jumps rather than changing the attribute by small amounts each period. The presence of infrequent but large jumps fits naturally with a model where there are private information shocks to the costs of investment, in which case our method may be useful. We now turn to empirical papers that compute this type of model.

4.2 Shipbuilding subsidies in Kalouptsidi (2018)

Kalouptsidi studies the market for ship building in a dynamic oligopoly model. Each period in the Kalouptsidi model, each shipyard i decides on a_i^t , which indicates how many ships to build. The shipyard's state, s_i^t , includes its backlog of orders and characteristics. The shipyard can sell its ships at a price, VE , which reflects the demand for new ships by ship owners, who are ship buyers. This price depends only on the exogenous portion of the state, s_0^t , so they write $VE(s_0^t)$. Production cost has a deterministic component that is a function

⁸A complication of having permanent entry and exit is that there are a potentially infinite number of players, though Pakes and McGuire (1994) cap the number of active players at any period. This requires notational changes but does not fundamentally change the nature of the game.

of the amount produced and the shipyard’s backlog, $\bar{c}(a_i^t, s_i^t, s_0^t)$.⁹ Production cost also has a stochastic normally-distributed component with standard deviation σ^t . The period payoff is thus given by:

$$\underbrace{VE(s_0^t) \times a_i^t - \bar{c}(a_i^t, s_i^t, s_0^t)}_{\pi_i(a_i^t, s^t)} - \underbrace{a_i^t \sigma_i^t}_{c_i(a_i^t, s^t)} \times \varepsilon_i^t, \quad (11)$$

where we have indicated the mapping to our notation with under braces.

Translating to our framework, the first two terms make up the deterministic component $\pi_i(a_i^t, s^t)$ and the third term the random component $c_i(a_i^t, s^t) \times \varepsilon_i^t$. Kalouptsidi (2018, p. 1123) shows that convexity of the cost function $\bar{c}(a_i^t, s_i^t, s_0^t)$ is sufficient for all investment levels being chosen with positive probability. The probability of a given investment level is given by the mass of ε_i^t falling between two neighboring difference in conditional choice values, corresponding to our cutoffs.

Thus, the Kalouptsidi (2018) result, applied to our context, essentially boils down to an assumption that every choice—in terms of the number of ships to build—lies in the discrete concave hull of choice-specific value functions at each state. Our results complements this paper by showing how to compute optimal policies for general payoff functions which do not necessarily induce discrete concave choice-specific values.

4.3 Digital movie adoption in Caoui (2023)

Caoui studies adoption of digital movie screens, also in a dynamic oligopoly setting. In Caoui, a movie theater’s action a_i^t represents the number of digital movie screens it adopts. The industry state s_0^t includes the mean price of installing digital screens, p^t , and the aggregate share of digital screens, h^t , while s_i^t is the number of screens the theater has previously adopted. There is also a normally distributed shock to the price of installing a digital movie screen, ε_i^t . Mean profits are a function of these states, $\bar{\pi}(s_i^t, h^t)$.

⁹Kalouptsidi (2018) refers to this term as c , but we use \bar{c} because we use c for the random component of the state. Similarly, we adjust the notation of other papers below.

The period payoff is given by:

$$\underbrace{\bar{\pi}_i(s_i^t, h^t) - a_i^t \times p^t}_{\pi_i(a_i^t, s^t)} - \underbrace{a_i^t}_{c_i(a_i^t, s^t)} \times \varepsilon_i^t. \quad (12)$$

Once again, we can divide profits into the deterministic component and a stochastic component that multiplies the unobservable term.

Caoui (2023, p. 610) assumes that the period payoff function satisfies “decreasing differences” in (a_i^t, ε_i^t) to ensure that the optimal investment choice a_i^t is monotone in ε_i^t . This assumption is similar to the choice-specific value function being discrete concave across options. In general, the choice-specific value function need not be discrete concave and consequently, not all actions are chosen with positive probability. Our algorithm in Proposition 2 allows the researcher to find the set of actions chosen with positive probability.

4.4 Energy transitions in Gowrisankaran et al. (2024)

Gowrisankaran et al. model a regulated monopoly utility faced with an energy transition. Each three-year period, the utility first decides how much coal capacity to retire and then how much combined-cycle natural gas (CCNG) capacity to add. The exogenous state variable s_0 is the market price for natural gas. The firm’s own state s_1 is formed from its current coal and CCNG capacities. The mean cost of retiring coal capacity is quadratic in the amount of capacity retired, while the mean cost of adding CCNG capacity is quadratic in the amount of capacity added. In both cases, there are also linear *i.i.d.* marginal cost shocks in c_1 . For either generation source, we can write the action a_1^t as the next period’s capacity level of that source. The cost function for retiring coal capacity can then be written as:

$$\underbrace{\bar{\pi}_1(s_1^t, s_0^t) - \delta_0 \mathbb{1}\{a_1^t \neq 0\} - a_1^t (\delta_1 + a_1^t \delta_2)}_{\pi_1(a_1^t, s^t)} - \underbrace{a_1^t \sigma}_{c_1(a_1^t, s^t)} \times \varepsilon_1^t. \quad (13)$$

Equation (13) considers a single investment decision, but in Gowrisankaran et al., the regulated monopoly makes two investment/retirement decisions each period. Because the two decisions are made in sequence, we can treat them as being made in separate periods

within the context of our model. The paper models 10 discrete levels of coal retirement and CCNG investment, and found that increasing the number of levels resulted in very similar structural parameter estimates.

Gowrisankaran et al. (2024) use the methods developed in this paper. In particular, they estimate the model with a full solution nested fixed point generalized method of moments (GMM) approach. For each candidate parameter value, they solve for the distribution of retirement/investment outcomes using the algorithm in Proposition 2. They then match moments of the state-contingent investment outcomes to the data. Their moments include the probabilities of retirement and investment, the retirement/investment amounts and their squares conditional on non-zero levels, and the standard deviations of these amounts. In Gowrisankaran et al. (2024), the choice-specific value function was not discrete concave at the estimated parameters and hence only a subset of retirement/investment levels were chosen with positive probability.

Gowrisankaran et al. (2024) is the special (monopoly) case of an oligopoly framework where each firm is faced with stochastic costs. In the more general case with N firms, the publicly observable state vector s^t consists of $N + 1$ elements, denoting the demand state s_0^t , and firms' capacities $(s_1^t, s_2^t, \dots, s_N^t)$. Each period, firms engage in Bertrand or Cournot competition given their capacity levels, and earn profits, $\bar{\pi}(s^t)$ based on these decisions. Firms face a deterministic linear quadratic asymmetric mean investment cost function. The stochastic part of the investment cost is proportional to investment, $a_i^t - s_i^t$. To model the specificity of capital, both the deterministic and stochastic parts are asymmetric around 0. Choosing the lowest level of capacity $a_i^t = 0$ corresponds to exit, in which case the firm would receive a scrap value. A firm that is already at capacity $s_i^t = 0$ would need to pay an entry cost to build capacity.

Firm i 's period payoff becomes:

$$\begin{aligned}
\bar{\pi}_i(s^t) &= \mathbb{1}\{a_i^t > s_i^t\} (\delta_1 + \delta_2(a_i^t - s_i^t) + \delta_3(a_i^t - s_i^t)^2 + \mathbb{1}\{s_i^t = 0\}\chi) \\
&- \mathbb{1}\{a_i^t < s_i^t\} (\delta_4 + \delta_5(a_i^t - s_i^t) + \delta_6(a_i^t - s_i^t)^2 - \mathbb{1}\{a_i^t = 0\}\phi) \\
&- [\mathbb{1}\{a_i^t < s_i^t\}(a_i^t - s_i^t)\sigma_1 + \mathbb{1}\{a_i^t > s_i^t\}(a_i^t - s_i^t)\sigma_2] \times \varepsilon_i^t.
\end{aligned} \tag{14}$$

In equation (14), the first two lines make up the deterministic component of period payoff $\pi_i(a_i^t, s^t)$. The first line indicates profits from the product market minus the deterministic cost of positive investment, and the second line is the deterministic part of negative investment. The parameters χ and ϕ denote entry cost and scrap values respectively. The third line contains the random component of investment cost $c_i(a_i^t, s^t) \times \varepsilon_i^t$. Different parameters σ_1 and σ_2 allow for asymmetry in the cost of positive and negative investment also in the random component.

This oligopoly model is similar to the Ryan (2012) and Fowle et al. (2016) model of the cement industry. However, these papers model unobservables to the fixed cost of investment rather than linear shocks that are proportional to the quantity of investment. Thus, in these papers, the sign of investment in any period—negative, positive, or zero—is stochastic, but the level of investment conditional on the sign, the level of investment is deterministic.

5 Conclusion

This paper develops new methods to compute a class of dynamic oligopoly models. We consider models where firms can invest to build or retire capacity or other attributes with many fixed values, that can approximate a continuous distribution. In our model, firms can adjust their capacity as much as they want in any period. The desired investment quantities are limited by quadratic costs of investment and capital specificity. Firms bear a private information component to the cost of investment. The private information component is a shock to the marginal cost of investment, so it is proportional to the amount invested.

We characterize the optimal investment policy for this type of model. Specifically, we show that the optimal policy is a set of investment options that are chosen with positive probability for some value of the private information shock. These actions are chosen with positive probability if and only if they are in the discrete analog of the concave hull of the choice specific value function relative to other actions. We develop a computationally efficient method to calculate the probability of choosing each action. Our method requires fewer than $2K$ iterations on a comparison process per state and firm, where K is the number of actions,

i.e. potential investment levels.

In contrast to our method, we show that equilibrium may not exist for approximate models that simulate the private information shocks and use the simulation draws to compute the probability of different investment levels. We provide a simple example of a 2×2 game where equilibrium does not exist with a finite number of simulation draws.

Our methods may be useful in analyzing dynamic oligopoly models with investment in fixed attributes such as capacity. They may allow researchers to extend investment models used in industrial organization to incorporate private information shocks on the marginal cost of investment and with fewer assumptions that are currently made. They may also allow researchers the ability to extend macroeconomic models of investment to oligopoly settings.

References

- Amir, R. (2005). Supermodularity and complementarity in economics: An elementary survey. *Southern Economic Journal*, 71(3):636–660.
- Benkard, C. L. (2004). A dynamic analysis of the market for wide-bodied commercial aircraft. *The Review of Economic Studies*, 71(3):581–611.
- Besanko, D. and Doraszelski, U. (2004). Capacity dynamics and endogenous asymmetries in firm size. *RAND Journal of Economics*, pages 23–49.
- Besanko, D., Doraszelski, U., Kryukov, Y., and Satterthwaite, M. (2010). Learning-by-doing, organizational forgetting, and industry dynamics. *Econometrica*, 78(2):453–508.
- Caoui, E. H. (2023). Estimating the costs of standardization: Evidence from the movie industry. *The Review of Economic Studies*, 90(2):597–633.
- Collard-Wexler, A. (2013). Demand fluctuations in the ready-mix concrete industry. *Econometrica*, 81(3):1003–1037.

- Cooper, R. W. and Haltiwanger, J. C. (2006). On the nature of capital adjustment costs. *The Review of Economic Studies*, 73(3):611–633.
- Doraszelski, U. and Markovich, S. (2007). Advertising dynamics and competitive advantage. *The RAND Journal of Economics*, 38(3):557–592.
- Doraszelski, U. and Satterthwaite, M. (2010). Computable Markov-perfect industry dynamics. *The RAND Journal of Economics*, 41(2):215–243.
- Ericson, R. and Pakes, A. (1995). Markov-perfect industry dynamics: A framework for empirical work. *The Review of Economic Studies*, 62(1):53–82.
- Fowlie, M., Reguant, M., and Ryan, S. P. (2016). Market-based emissions regulation and industry dynamics. *Journal of Political Economy*, 124(1):249–302.
- Gowrisankaran, G. (1995). A dynamic analysis of mergers. Doctoral Dissertation, Yale University.
- Gowrisankaran, G. (1999). A dynamic model of endogenous horizontal mergers. *The RAND Journal of Economics*, pages 56–83.
- Gowrisankaran, G., Langer, A., and Reguant, M. (2024). Energy transitions in regulated markets. Working Paper.
- Gowrisankaran, G. and Town, R. J. (1997). Dynamic equilibrium in the hospital industry. *Journal of Economics and Management Strategy*, 6:45–74.
- Hotz, V. J. and Miller, R. (1993). Conditional choice probabilities and the estimation of dynamic models. *Review of Economic Studies*, 60:497–529.
- Kalouptsi, M. (2018). Detection and impact of industrial subsidies: The case of chinese shipbuilding. *The Review of Economic Studies*, 85(2):1111–1158.
- Khan, A. and Thomas, J. K. (2008). Idiosyncratic shocks and the role of nonconvexities in plant and aggregate investment dynamics. *Econometrica*, 76(2):395–436.

- Maskin, E. and Tirole, J. (1988). A theory of dynamic oligopoly 1: Overview and quantity competition. *Econometrica*, 56:549–599.
- Mermelstein, B., Nocke, V., Satterthwaite, M. A., and Whinston, M. D. (2020). Internal versus external growth in industries with scale economies: A computational model of optimal merger policy. *Journal of Political Economy*, 128(1):301–341.
- Pakes, A., Gowrisankaran, G., and McGuire, P. (1992). Code for implementing the pakes-mcguire algorithm for computing Markov perfect equilibrium. Working Paper, Yale University.
- Pakes, A. and McGuire, P. (1994). Computing Markov-perfect Nash equilibria: Numerical implications of a dynamic differentiated product model. *RAND Journal of Economics*, 25(4):555–589.
- Rust, J. (1987). Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica*, 55:999–1033.
- Rust, J. and Rothwell, G. (1995). Optimal response to a shift in regulatory regime: the case of the US nuclear power industry. *Journal of Applied Econometrics*, 10(S1):S75–S118.
- Ryan, S. P. (2012). The costs of environmental regulation in a concentrated industry. *Econometrica*, 80(3):1019–1061.