

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, Revised March 2025

We thank Patricio Hernandez and Filippo Paternollo for excellent research assistance. We thank Hassan Afrouzi, El Hadi Caoui, Allan Collard-Wexler, Jean-François Houde, Mitsuru Igami, Fedor Iskhakov, Paul Koh, anonymous referees, and seminar participants at the University of Arizona, CREST-INSEE, Università Cattolica del Sacro Cuore, and the University of Vienna for helpful comments. Gowrisankaran acknowledges financial support from the National Science Foundation through Grant SES-2149335. 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.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, Revised March 2025
JEL No. C63, C73, L11, L13

ABSTRACT

We analyze a class of dynamic models that has several recent applications, where each period, each firm receives a private shock to the marginal cost of investment and chooses among many ordered capacity levels. Simulation methods to compute these models can result in non-existence of pure strategy equilibrium, while including multinomial shocks leads to counterintuitive predictions. We provide a computationally fast method to calculate optimal investment probabilities given value functions that iteratively finds investment levels chosen with positive probability and cutoffs of private information shocks across choices. Our method is useful to practitioners seeking to estimate models with continuous firm choices.

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 and empirical literature on dynamic models, often with many endogenously heterogeneous firms interacting (Ericson and Pakes, 1995). One of the most common ways that firms can differentiate themselves is through investment decisions that affect their future state variable, which we can broadly think of as capacity. These investment decisions generally have many ordered choices or are continuous, with firms often facing convex costs that limit how much capacity they will add in each year. Many empirical papers have studied decisions that fit within this framework, including nuclear plant ramping (Rust and Rothwell, 1995), cement capacity investment (Ryan, 2012), concrete plant entry and growth (Collard-Wexler, 2013), shipyard queues (Kalouptsidei, 2018), individual asset level changes (Chatterjee et al., 2023), digital movie screen adoption (Caoui, 2023), electricity generation capacity investments (Gowrisankaran et al., 2024), and rooftop solar installations (Snashall-Woodhams, 2024). These models have substantially added to our understanding of industry dynamics, including by providing evidence on the impact of supply chain hold-ups, carbon taxes, and technology adoption, among other policy-relevant economic phenomena.

This paper develops methods to study capacity dynamics in models with many ordered choices and potentially strategic interactions between firms. While we formally model the case of firms that are faced with a large number of ordered choices, our methods can also be used to analyze the limiting case of dynamic agents faced with a continuous capacity choice each period. Hence, our paper may be useful in analyzing the types of research questions discussed above.

We consider a model where each period, each firm obtains a private information shock to the marginal cost of investment that it learns before it makes its investment decision. We propose a computationally fast algorithm to compute the best responses of firms in these settings. The input to our algorithm is a set of choice-specific value functions, one for each of the large number of choices, where the choice-specific values do not include the unobservable term. The output is the probability of each investment level and the resulting

ex-ante value function. We suggest nesting our algorithm within the best response iteration method used to compute equilibria for the canonical dynamic oligopoly models (Pakes and McGuire, 1994; Ericson and Pakes, 1995). This method involves repeatedly solving firms' Bellman equations and beliefs about rivals' decisions. Our technique can also be used to solve for the optimal best responses in a conditional choice probability (CCP) estimator for models with continuous choices (Granja and Forte, 2024), which might provide more desirable econometric properties relative to inequality-based estimators that do not solve for best responses (Bajari et al., 2007).

A fundamental component of our model is the privately observed shock to the marginal cost of investment, which might be caused by random variation in the price of an input, for instance. In the context of dynamic oligopoly models, it is useful to include such shocks for at least two reasons. First, private information shocks are key to establishing existence of pure strategy Markov Perfect Equilibrium (Gowrisankaran, 1995; Doraszelski and Satterthwaite, 2010). Without private information shocks, the ex-ante probability of a firm choosing a given investment level may be discontinuous in its choice-specific value function, which will then result in a discontinuity in its rivals' value functions in their best responses to the firm's own value function. This prevents the application of Brouwer's fixed-point theorem, which could otherwise generally be used to show existence of pure strategy equilibrium.

Second, private information shocks generate *ex-ante* stochasticity in outcomes such as investment, which is important as a basis for empirical work, for both single-agent and dynamic game models. Specifically, as has been noted since at least Rust (1987), models with unobservable shocks are necessary to fit inherent randomness in the data and generate a well-defined likelihood.

An alternative to our model with unobservable shocks to the marginal cost of investment would be a model with multinomial cost shocks. In this case, the underlying investment decision would include an *i.i.d.* residual that is specific to each capacity choice and independent across choices.¹ This model would address the issues of existence of equilibrium and also gen-

¹For 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. More recent papers that include logit shocks include Collard-Wexler (2013), Chatterjee et al.

erate a well-defined likelihood. However, in this type of model, the number of shocks increases with more choices and hence a finer grid, implying that the number of choices in the approximation will have a big impact on the outcomes. We substantiate this point with simulation evidence that shows that the mean and standard deviation of investment and *ex-ante* values all fail to converge as the number of grid points increases. This lack of convergence suggests that the multinomial model may result in inconsistent structural parameter estimates if this model is used in estimation. This model also has the undesirable property that the shock to a large positive investment is independent of a shock to a medium-large positive investment, which may result in implausible counterfactual substitution patterns. Finally, this model also cannot fit the data in many contexts. For instance, manufacturing industries generally have lumpy investment, suggesting high fixed costs (Cooper and Haltiwanger, 2006). With substantial fixed costs, our simulation results show that investment increases as the number of choices increases. This occurs because every choice except one—no investment—involves paying the fixed cost and hence the proportion of choices with positive investment increases as the number of choices increases. Thus, in many contexts where choices are ordered, a model with unobservable shocks to the marginal cost of investment may fit the economics of the problem more closely than a model with multinomial shocks.²

Despite the importance of modeling dynamic oligopoly settings with continuous choices and the attractiveness of linear shocks to the marginal cost of investment in many of these settings, there has been no established framework to model continuous choices in a dynamic setting with unobservables. In our view, this is partly because researchers currently have a limited ability to compute best responses for this setting. The point of the paper is to address this limitation. One approach to computing best responses is to simulate the marginal cost shock, compute the optimal investment level for each shock draw, and numerically integrate across draws to approximate the probability of each investment level (Caoui, 2023).³ As we

(2023), and Snashall-Woodhams (2024).

²Other papers model a shock to the fixed cost of investment (Ryan, 2012), but this does not generate variation in outcomes conditional on a positive investment level and hence does not necessarily result in a continuous mapping from values and policies to best response values and policies.

³Simulation for single-agent continuous choice models is also common in macroeconomics (Khan and Thomas, 2008).

illustrate with an example, using simulation draws in this way does not result in a continuous mapping from a given set of values and policies to the best response values and policies. Even with large numbers of simulation draws, a pure strategy Markov Perfect Equilibrium may not exist for the approximate model with simulation draws, even if it exists for the limiting model with a continuous shock distribution.

We present, and build on, an alternative approach to computing best responses, which is to calculate the optimal cutoffs of the private information shock between different choices. Without any further simplification and with K potential investment choices, this would require comparing each investment choice against each other investment choice, for a total of $K(K - 1)/2$ comparisons, a number that grows quadratically in K . A simplification would involve assuming that each choice is chosen with positive probability, leveraging assumptions on the period payoff functions (Kalouptsi, 2018). Because the choices of investment are monotonic in the unobserved cost shock (Bajari et al., 2007), in this case it would then be sufficient to compute cutoffs for neighboring capacity choices, requiring only $K - 1$ comparisons. However, as verified by our simulation evidence, the assumption that all choices have positive probability may not hold when the choice-specific value functions are the result of complex optimization under uncertainty and interactions across firms in oligopoly models.

In contrast, we start by deriving the discrete analog of a concave hull for the choice-specific value functions. We show that an investment choice will be chosen with positive probability if and only if it lies in this discrete concave hull. We then develop an algorithm to find the discrete concave hull that requires at most $2K - 3$ comparisons. The idea is to iteratively compare an investment choice against its neighbor, and mark an investment level as being chosen with zero probability if the cutoff is in the wrong direction from the previous cutoff. In this instance, the algorithm works backward, potentially eliminating earlier choices that also have a cutoff in the wrong direction. The output of this algorithm is the set of investment choices that are chosen with positive probability, the probability of choosing each choice, and the cutoff values of shocks between choices. The computational time for our algorithm to derive a set of best response probabilities is linear in the number of investment

choices. The algorithm to compute these probabilities is less than 100 lines of code.⁴

We believe that our method is useful to practitioners wanting to empirically model continuous firm choices such as investment, which is true in many industrial organization applications. In these cases, our approach is valuable because unobservable shocks are important, shocks to the marginal cost of investment may be the most economically intuitive, and choice-specific value functions may not lie in a discrete analog of a concave hull.

Relation to literature. Our model builds on a literature on computing and estimating dynamic oligopoly models. These models have been used to theoretically analyze many important economic issues, including the evolution of market structure (Pakes and McGuire, 1994), the long-run impact of horizontal mergers and antitrust policy (Gowrisankaran, 1999; Mermelstein et al., 2020), the substantial and persistent firm sizes differences that are observed in most industries (Besanko and Doraszelski, 2004), and the role of learning-by-doing and forgetting and advertising in industry dynamics (Doraszelski and Markovich, 2007; Besanko et al., 2010a). This literature has also been used as a framework for empirical work where researchers first estimate the parameters of a dynamic model and then evaluate policy counterfactuals. This approach has been used to study the dynamics of hospital quality (Gowrisankaran and Town, 1997), aircraft costs (Benkard, 2004), brand equity for consumer goods (Borkovsky et al., 2017), and many other sectors.

A number of other papers seek to facilitate the computation of dynamic models with continuous choices. In the macroeconomics literature, the endogenous gridpoints method (Carroll, 2006) considers agents faced with a continuous decision (e.g., amount to consume versus save) where the agent’s period maximization decision has a closed form first-order condition. The method consists of starting with an exogenous grid of assets in some terminal period and then finding the (endogenous) asset levels in previous periods that lead to the later grid points by inverting a first-order condition. This method has been extended to consider discrete-continuous settings (Iskhakov et al., 2017). Our model differs from these models in that the agents in our model receive an unobservable shock prior to making their continuous

⁴We provide Python code at <https://github.com/patohdzs/gsd-capacity-investment>. This code includes a function that implements our algorithm, called `compute_optimal_cutoff`, and an example with 20 choices that calls this function.

decision. While this no longer allows for a one-to-one mapping between grid points across periods, it is useful in ensuring existence for dynamic oligopoly models and also in deriving models that can fit investment data. Finally, other papers extend the estimation of dynamic oligopoly papers in a number of ways, leveraging assumptions that the set of choices belongs to the discrete concave hull (Srisuma, 2013; Aradillas-López and Gandhi, 2016; Koh, 2022). Our paper is complementary to these papers in showing how to compute the discrete analog of the concave hull in a computationally fast way.

The remainder of this paper is divided as follows. Section 2 describes our model. Section 3 proposes computational methods, including our algorithm. Section 4 discusses applications from the literature. Section 5 provides simulation evidence comparing our model with linear shocks to the marginal cost of investment to a model with multinomial *i.i.d.* shocks. Section 6 concludes.

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 value 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 $i \in \mathcal{N}$. Each period, firms decide whether to invest and update their capacity level, choosing from a set of K potential capacity levels.

We note two general points regarding our model. First, although our model formally considers the case of K capacity levels, the K levels may also be a discrete approximation to a continuous investment decision. Second, while we define the state variable as capacity for ease of exposition, as we discussed in Section 1, our notion of capacity is quite general and applies to a variety of empirical settings, ranging from solar system size to asset levels held.

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. For use in describing the applications discussed in Section 4, we partition the state into endogenous and exogenous parts. We let s_i^t denote the portion of s^t that is affected by the actions of firm i and 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.⁵ The strict monotonicity assumption is equivalent to full support on \mathfrak{R} . Continuity of F allows us to ensure existence of pure strategy equilibria.⁶ The strict monotonicity and continuity assumptions together ensure that multiple investment outcomes are chosen with positive probability *ex-ante*. Finally, we assume that $\mathbb{E}[\varepsilon_i^t | \varepsilon_i^t \geq \varepsilon] < \infty$ for all ε and i , 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 a_i^t . We assume that a_i^t is chosen from the action set $A_i = (\alpha^1, \dots, \alpha^K)$; i.e. A_i consists of K unique real numbers. We use $o(\cdot)$ to denote the cardinality of sets, so $o(A_i) = K$. For ease of later analysis and without loss of generality, we impose an increasing order on A_i : let $\alpha^1 < \alpha^2 < \dots < \alpha^K$. Similarly, we impose an increasing order on the elements of any subset of A_i . 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) .

⁵Rust (1987) provides a discussion of the independence assumption in dynamic discrete choice models.

⁶See the discussion in Gowrisankaran (1995) or Doraszelski and Satterthwaite (2010).

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 on $S \times A_i$ given by:

$$\pi_i(a_i^t, s^t) - \tilde{c}_i(a_i^t, s^t) \times \varepsilon_i^t, \quad (1)$$

where $\tilde{c}_i(a_i^t, s^t)$ is the component of investment cost that multiplies the unobservable shock. Given that we are interested in linear marginal cost shocks, a typical specification would be $\tilde{c}_i(a_i^t, s^t) \equiv a_i^t - s_i^t$. We let $\tilde{c}_i(a_i, s)$ be strictly increasing in its first argument a_i , i.e., the firm's stochastic cost of capacity is increasing in the level of capacity chosen.⁷ The deterministic term $\pi_i(a_i, s)$ depends on the capacity choice for next period a_i^t and the current state of the industry s^t .⁸ Conceptually, $\pi_i(a_i, s)$ includes the profits from selling the product and the deterministic part of investment costs. Also for use in describing applications, we can separate $\pi_i(a_i, s)$ into these two components: $\pi_i(a_i, s) \equiv \bar{\pi}_i(a_i, s) - \bar{c}_i(a_i, s)$. A natural specification for $\bar{c}_i(a_i, s)$ would be a quadratic total cost function with three parameters, on fixed costs, quantity, and quantity squared (Ryan, 2012). We assume that both π_i and \tilde{c}_i are

⁷Mathematically, the monotonicity part of this assumption is without loss of generality, since we could simply reorder actions so that this assumption holds. However, our model can also be used as a finite approximation of a specification with continuous investment decisions along an interval, for which actions have a natural cardinal interpretation that cannot be altered.

⁸For ease of notation, we let period profits depend only on the firm's own actions given the state. However, it is easy to extend our model to allow period profits to also depend on its rivals' expected actions. This is useful in contexts such as learning-by-doing, where a firm's dynamic action a_i is its quantity choice and hence impacts its rivals' expected current profits (Benkard, 2004).

bounded: $|\pi_i(a_i, s)|, |\tilde{c}_i(a_i, s)| < \infty$ 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 , this discounted sum is given by:

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

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 shocks are 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 the realization of a specific action profile at each publicly observable state— $\omega_i(a|s)$, which is the product of beliefs about the probability of its own action multiplied by beliefs about the probability of rival players' actions, $\omega_i(a|s) \equiv \omega_i(a_i|s) \times \omega_i(a_{-i}|s)$.

Value function. Following the literature on dynamic games in industrial organization,⁹ we define a Bellman equation for firm i given any set of beliefs for the probabilities of actions, ω_i . Focusing on the ex-ante value function, i.e. before the private shock ε_i is realized, we

⁹See for example Aguirregabiria et al. (2021).

obtain:

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

Here, \mathbb{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 expoit 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 function*:

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

as firm i 's value net of the random component of cost $\tilde{c}_i(a_i, s) \times \varepsilon_i$ when it chooses action k . It is optimal for it to choose action k under beliefs ω_i whenever

$$v_i^k(s|\omega_i) - \tilde{c}_i(\alpha^k, s) \times \varepsilon_i \geq v_i^\ell(s|\omega_i) - \tilde{c}_i(\alpha^\ell, s) \times \varepsilon_i, \forall \ell \neq k. \tag{5}$$

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:

$$a_i(\varepsilon_i, s) = \operatorname{argmax}_{\alpha^k \in A} \left\{ v_i^k(s|\omega_i) - \tilde{c}_i(\alpha^k, s) \times \varepsilon_i \right\}. \tag{6}$$

The probability that firm i chooses action k in state s thus given by

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

This relationship holds for all firms $i \in \mathcal{N}$, states s , and actions k . 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 $L \times N \times (K - 1)$ equations, which we can write compactly in vector notation as:

$$p = \psi(\omega), \tag{7}$$

where p denotes the $L \times N \times (K - 1)$ -dimensional vector of choice probabilities and ω the $L \times N \times (K - 1)$ -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 $a_i(\varepsilon_i, s)$ 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 an $L \times N \times (K - 1)$ -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 with a restriction to symmetric strategies in the presence of symmetric payoffs.

3 Computational algorithms

In order to compute a dynamic equilibrium of the model, we propose to 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, we propose to start with beliefs of rivals' choice probabilities, ω_i , and the corresponding ex-ante value function, $V_i(s|\omega_i)$, for each $i \in \mathcal{N}$. The algorithm updates these matrices as follows. For every state s and every firm 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.

While this algorithm is standard in the literature, our innovation is in the updating of the ex-ante value function $V_i(s|\omega_i)$. To update $V_i(s|\omega_i)$, one needs to integrate over the ε draws, evaluate the chosen action for each draw, and calculate the resulting value. We propose new methods to perform this integration quickly and without using simulation.¹⁰

3.1 Simulation-based computation and nonexistence of equilibria

Before turning to our method, we first discuss issues with the standard approach, which is to calculate this integral by simulation. With this 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.

¹⁰While our method improves the speed of each step of the successive approximations algorithm, it does not address the computational problem of a curse of dimensionality from large state spaces (Rust, 1997).

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 solved via simulation, the approximated 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.¹¹

In an earlier working paper that estimated dynamic oligopoly equilibria for hospital investment models (Gowrisankaran et al., 2010), we found that simulation methods did not converge. This working paper led to the development of the current paper. We now show the potential for non-existence with two (simpler) 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 1 (i.e., $a_1 = \alpha^1$) 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 2.¹² For simplicity, we consider uniformly distributed shocks with $\varepsilon_i \sim U(-0.5, 0.5)$.¹³ We give the payoff matrix for the modified game below:

¹¹Traditional quadrature and importance sampling methods would still result in reaction functions that are discontinuous in the values of other participants, and hence not solve this issue.

¹²This matching pennies model corresponds to a capacity adjustment model where there is a stochastic cost when the firm needs to adjust capacity to 2.

¹³Unlike our model in Section 2, which assumes full support of ε_i over \mathfrak{R} , the support of the uniform distribution is bounded. While we choose this density for ease of exposition, it is possible to extend this example to a density with full support.

		Player 2	
		$a_2 = \alpha^1$	$a_2 = \alpha^2$
Player 1	$a_1 = \alpha^1$	$(-.75, 1)$	$(1.25, -1 - \varepsilon_2)$
	$a_1 = \alpha^2$	$(1 - \varepsilon_1, -1)$	$(-1 - \varepsilon_1, 1 - \varepsilon_2)$

Let p_i be Player i 's probability of choosing action 1. Player 1 chooses action 1 if and only if $\varepsilon_1 > 4p_2 - 2.25$. Player 2 chooses action 1 if and only if $\varepsilon_2 > 2 - 4p_1$. The modified game has a unique pure strategy equilibrium. Player 1 plays action 1 if and only if $\varepsilon_1 > -1/68$. Similarly, Player 2 plays action 1 if and only if $\varepsilon_2 > -1/17$. First suppose 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 2, when drawing $\varepsilon_i = -0.3$ and action 1 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 1 if and only if her draw was larger than $4p_2 - 2.25 = -0.25$, she would never choose action 2, which implies that she chooses the same action, action 1, 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 a variant of our 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 $\pi(1)$, if two firms remain active, they each earn duopoly profits $\pi(2)$.¹⁴ Firms that have exited earn profits of zero. Assume that $\pi(2) < 0 < \pi(1)$, so that, without 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. $\tilde{c}(\alpha^k) = \alpha^k$, 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:

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

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 would 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 following four sets of values:

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

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 one's 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

¹⁴For ease of notation, we suppress the dependence of π on a_i or i .

not exist for some values of the draws for the approximated model with discretized shocks.

3.2 Characterization of action choice probabilities

Having shown the issues with a simulation approach, we now turn to methods that find the exact action choice probabilities, which includes our method. This subsection starts by showing which choices k are chosen with positive probability, and the probability that each capacity level is chosen given that it is chosen with positive probability. The results in this subsection serve as a precursor to our main innovation, which is a fast algorithm to compute these probabilities.

In the remainder of this section, we condition on a state s and a firm i and the accompanying choice-specific values for this state. Since we now consider only the decision of a particular state and firm, we use $a(\varepsilon)$, $\tilde{c}(\alpha^k)$, and v^k to refer to $a_i(s, \varepsilon)$, $\tilde{c}_i(\alpha^k, s)$, and $v_i^k(s|\omega_i)$ respectively, to ease notation.

We start with the following lemma:

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

Proof Because the period payoffs exhibit decreasing differences in (a, ε) , the result follows from Topkis' Theorem, but we prove it directly.¹⁵ Consider two actions j and k , with $j < k$:

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

where the third line uses the fact that $\tilde{c}(\cdot)$ is increasing, implying that $\tilde{c}(\alpha^k) - \tilde{c}(\alpha^j)$ 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:

¹⁵Bajari et al. (2007) also use decreasing differences assumption to show monotonicity, referring to it as “Monotone Choice.”

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

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

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 action $k \iff \varepsilon > \bar{\varepsilon}(j, k)$.

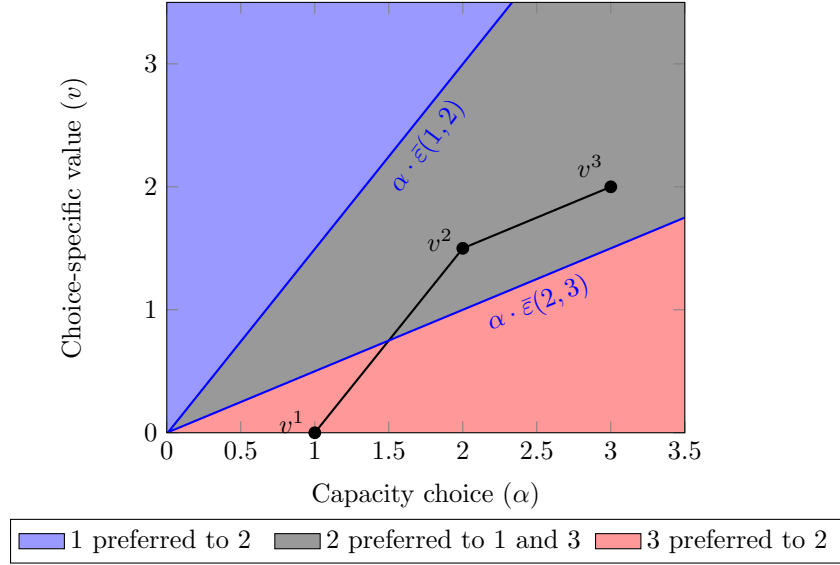
Proof This result follows directly from the proof of Lemma 1. ■

In some cases, one can directly use the ε cutoffs to define the probabilities of each action. However, in other cases, some actions are never chosen and in these instances, one cannot directly use the ε cutoffs. Before turning to the general results, we illustrate such zero probability actions with an example. Consider the case of three possible actions, $K = 3$, and linear capacity costs with slope 1, so $\tilde{c}(\alpha^k) = k$. The 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, the 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 $0.5 < \varepsilon \leq 1.5$, and $a(\varepsilon) = 3$ when $\varepsilon \leq 0.5$. Consequently, all three actions are chosen with positive probability.¹⁶ 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

¹⁶Note that in Section 2.2 we adopted the convention that at the cutoff the firm chooses the larger action.

Figure 1: All actions chosen with positive probability



not coincidentally, using Lemma 2, one can verify that $a(\varepsilon) = 1$ when $\varepsilon > 1$, and $a(\varepsilon) = 3$ when $\varepsilon < 1$. Consequently, $a(\varepsilon) = 2$ is never optimal.

Overall then, Lemma 2 and Figures 1 and 2 imply that the presence of decreasing differences in the period payoff functions—or equivalently in $\tilde{c}(a_i) \times \varepsilon$ —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.

To formalize this intuition, we first introduce a definition of the best ε cutoff for any option as being the outer-most of all the cutoffs in one direction:

Definition Best ε cutoffs for option k .

1. Define the *best lower ε cutoff* for option k in set \tilde{C} as $\bar{\varepsilon}^L(k, \tilde{C}) = \max\{-\infty, \max_{l>k, l \in \tilde{C}}\{\bar{\varepsilon}(k, l)\}\}$.
2. Define the *best upper ε cutoff* for option k in set \tilde{C} as $\bar{\varepsilon}^U(k, \tilde{C}) = \min\{\infty, \min_{j<k, j \in \tilde{C}}\{\bar{\varepsilon}(j, k)\}\}$.

Figure 2: Only two actions chosen with positive probability



We now offer our result that characterizes when actions are chosen with positive probability in terms of the ε cutoffs and the discrete concave hull:

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 $j < k < l$, it holds that $v^k > \lambda v^j + (1 - \lambda)v^l$, where $\lambda = \frac{\bar{c}(\alpha^k) - \bar{c}(\alpha^j)}{\bar{c}(\alpha^l) - \bar{c}(\alpha^j)}$.*

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 action k since it prefers action j to action 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 action k in this case since it prefers l to k . Thus, the firm picks action 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 $j < k < l$, $\bar{\varepsilon}(k, l) < \bar{\varepsilon}(j, k)$. Now consider $\varepsilon \in [\bar{\varepsilon}^L(k, A), \bar{\varepsilon}^U(k, A)]$, i.e. values of ε that lie above the best lower ε cutoff and below the best upper ε cutoff for every choice in A . For any such

ε , action k is preferred to action j , $\forall j < k$ and k is preferred to l , $\forall l > k$, and thus the firm chooses action k . Moreover, this interval is non-empty by assumption. Thus, action k is chosen with positive measure or probability.

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

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

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 as defined in the third statement of Proposition 1. The proposition also implicitly provides a way of computing which actions are chosen with positive probability and the accompanying probability of each action. We formalize:

Corollary 1. *Action j is chosen with positive probability if $\bar{\varepsilon}^L(j, A) < \bar{\varepsilon}^U(j, A)$. In this case, it is the weakly preferred option when $\varepsilon \in [\bar{\varepsilon}^L(j, A), \bar{\varepsilon}^U(j, A)]$. Finally, $Pr(a = \alpha^j) = Pr(\bar{\varepsilon}^L(j, A) \leq \varepsilon < \bar{\varepsilon}^U(j, A))$.*

Proof These results follows from the proof of Proposition 1. \blacksquare

Overall, Proposition 1 and Corollary 1 serve two purposes. First, they formalize a characterization that may not have generally been widely known. For instance, in the industrial organization literature, Kalouptsi (2018) and Caoui (2023) did not employ the results of this proposition, instead using a supply side assumption that guaranteed that all options are chosen with positive probability and simulation methods, respectively. In the macroeconomics literature, Cooper and Haltiwanger (2006) and Khan and Thomas (2008) use simulation methods to solve this type of problem.

Second, the results from Proposition 1 are necessary to prove our main result. Specifically, while Corollary 1 implicitly provides an algorithm for calculating the probability of each

action, 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 . Our main result, in Proposition 2 below, shows how to compute these probabilities with a number of calculations that grow linearly with K .

3.3 Exact and fast algorithm

This subsection presents the main innovation of this paper, which is an algorithm to update the ex-ante value function $V_i(s|\omega_i)$ that requires a number of calculations that grows only linearly with K . Specifically, conditional on choice-specific values for a state (s, i) , our algorithm allows us to find the exact probabilities of choosing each choice k and the accompanying ex-ante value function. Since our algorithm does not use simulation and is quick to compute, it is suitable to use nested within the dynamic game solution.¹⁷ 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 $C(A)$. 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:¹⁸
 - 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:

¹⁷Note that our algorithm only concerns the updating of the ex-ante value function $V_i(s, \omega_i)$, which can be used within a standard successive approximations method or a CCP estimator with continuous choices.

¹⁸We use \hat{C}_k to denote the k th element of \hat{C} .

- 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 Note first that because of the full support assumption for ε , actions 1 and K are always chosen with positive probability. Hence the algorithm can start with $k = 2$ and end with a comparison of the penultimate element to action K .

We prove the proposition in two steps. First, we show that $C(A) \subseteq \hat{C}$. Then we show that $\hat{C} \subseteq C(A)$.

Part 1: $k \in C(A) \Rightarrow k \in \hat{C}$. Consider any $k \in C(A)$. The only way it is possible that k is not in \hat{C} is if it were dropped by the algorithm from the candidate set, which occurs in Step 3. Step 3 drops k if, for some j, l , $j < k < l$, $\bar{\varepsilon}(j, k) \leq \bar{\varepsilon}(k, l)$. From Corollary 1, $\bar{\varepsilon}^U(k, A) > \bar{\varepsilon}^L(k, A)$, since $k \in C(A)$. But, by the definitions of the best ε cutoffs, $\bar{\varepsilon}(j, k) \geq \bar{\varepsilon}^U(k, A) > \bar{\varepsilon}^L(k, A) \geq \bar{\varepsilon}(k, l)$. Thus, element k could not have been dropped in the algorithm, implying that $k \in \hat{C}$.

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

We next show a second claim: for all $k \notin C(A)$, $\exists j, l \in C(A)$, $j < l$, such that $\bar{\varepsilon}(j, k) \leq \bar{\varepsilon}(k, l)$. Suppose, again by contradiction, that $\exists k \notin C(A)$ such that $\bar{\varepsilon}(j, k) > \bar{\varepsilon}(k, l), \forall j, l \in C(A)$. Then, consider $\varepsilon \in (\bar{\varepsilon}^L(k, C(A)), \bar{\varepsilon}^U(k, C(A)))$, i.e., values of ε for which choice k is the best option in the set $C(A)$. 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 Proposition 1, which shows that elements $k \notin C(A)$ are chosen with zero probability.

Now consider the possibility that $\exists k$ such that $\hat{C}_k \notin C(A)$. By the second claim, $\exists j', l' \in C(A)$, $j' < \hat{C}_k < l'$, such that $\bar{\varepsilon}(j', \hat{C}_k) \leq \bar{\varepsilon}(\hat{C}_k, l')$. By Part 1 of the proof, $j', l' \in \hat{C}$. Denote these elements \hat{C}_j and \hat{C}_l respectively. Then, by the first claim,

$$\bar{\varepsilon}(\hat{C}_{l-1}, \hat{C}_l) < \dots < \bar{\varepsilon}(\hat{C}_{k+1}, \hat{C}_k) < \bar{\varepsilon}(\hat{C}_{k-1}, \hat{C}_k) < \dots < \bar{\varepsilon}(\hat{C}_j, \hat{C}_{j+1}).$$

We now show a third claim: $\bar{\varepsilon}(\hat{C}_{k-1}, \hat{C}_k) \leq \bar{\varepsilon}(\hat{C}_j, \hat{C}_k) \leq \bar{\varepsilon}(\hat{C}_j, \hat{C}_{j+1})$. We prove this claim by contradiction. Suppose $\bar{\varepsilon}(\hat{C}_j, \hat{C}_k) < \bar{\varepsilon}(\hat{C}_{k-1}, \hat{C}_k)$. Then, for ε in a neighborhood immediately to the right of $\bar{\varepsilon}(\hat{C}_j, \hat{C}_k)$, \hat{C}_j is preferred to \hat{C}_k . But, at this point, \hat{C}_k is preferred to \hat{C}_{k-1} , which is preferred to \hat{C}_{k-2} etc. and is ultimately preferred to \hat{C}_j , by the first claim. Thus, this yields a contradiction. Similarly, if $\bar{\varepsilon}(\hat{C}_j, \hat{C}_k) > \bar{\varepsilon}(\hat{C}_j, \hat{C}_{j+1})$ there would be an equivalent contradiction.

We can also show that $\bar{\varepsilon}(\hat{C}_{l-1}, \hat{C}_l) \leq \bar{\varepsilon}(\hat{C}_k, \hat{C}_l) \leq \bar{\varepsilon}(\hat{C}_k, \hat{C}_{k+1})$ using analogous arguments to those we used to show the third claim.

Thus, $\bar{\varepsilon}(\hat{C}_k, l') = \bar{\varepsilon}(\hat{C}_k, \hat{C}_l) < \bar{\varepsilon}(\hat{C}_j, \hat{C}_k) = \bar{\varepsilon}(j', \hat{C}_k)$ which yields a contradiction to our assumption that $\bar{\varepsilon}(j', \hat{C}_k) \leq \bar{\varepsilon}(\hat{C}_k, l')$, showing that $\nexists k$ such that $\hat{C}_k \notin C(A)$. Thus, $C(A) = \hat{C}$.

■

The algorithm in Proposition 2 creates the set $C(A)$, while simultaneously defining the ranges of ε for which each action is chosen. To formalize the second part of the algorithm, we offer:

Corollary 2. *We can calculate the output of the Proposition 2 algorithm as follows:*

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

2. The probabilities of actions are:

(a) Action 1: $Pr(\varepsilon > \bar{\varepsilon}(\hat{C}_{o(\hat{C})-1}, \hat{C}_{o(\hat{C})}))$.

(b) Action k , $1 < k < K$: $Pr(\bar{\varepsilon}(\hat{C}_k, \hat{C}_{k+1}) \leq \varepsilon < \bar{\varepsilon}(\hat{C}_{k-1}, \hat{C}_k))$.

(c) Action K : $Pr(\varepsilon \leq \bar{\varepsilon}(\hat{C}_{o(\hat{C})-1}, \hat{C}_{o(\hat{C})}))$.

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

A practitioner may want to use the Proposition 2 algorithm to find $C(A)$ and then apply Corollary 2 to recover action probabilities and the *ex-ante* value function. This process can then be embedded within a CCP estimator or an equilibrium computation.

As we discussed above, a principal 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 3. *The algorithm defined by Proposition 2 computes at least $K - 1$ and at most $2K - 3$ ε cutoffs.*

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 speed of computation is important because the algorithm outlined in Proposition 2 will be performed many times for many states in order to solve for an MPE of a dynamic oligopoly model.

4 Applications From the Literature

We discuss a dynamic oligopoly literature that largely builds on the canonical Pakes and McGuire (1994) (and Ericson and Pakes, 1995) model. We thus first explain how our results relate to Pakes and McGuire (1994) and why our method is not useful in simplifying 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 choose future capacity levels with many ordered choices. 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 at period t indicates the quality level of each firm relative to the outside option. Each period t , firms invest a continuous amount $a_i^t \geq 0$ 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. At the same time, quality relative to the outside good decreases by 1 with some exogenously given probability.¹⁹ 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 and builds on it in three different ways. First, and most importantly, Pakes and McGuire consider the case where the state can vary over only one of a small number of levels from period to period.²⁰ The reason the state evolves slowly in Pakes and McGuire is that the original model concerns quality and it is natural to think of quality as evolving slowly. In contrast, capacity can change by a large amount from period to period empirically, albeit at a high cost. Our capacity investment model is designed for the case where there are many ordered choices and hence it is not directly useful to the case

¹⁹Extensions have modeled active disinvestment (Besanko et al., 2010b).

²⁰Pakes and McGuire allow quality to increase or decrease by at most one level, but empirical applications allow for quality to increase by two units per period (Borkovsky et al., 2017).

where there are only a small number of potential levels in the following period.

Second, in Pakes and McGuire, there is no privately observed cost shock. Instead, Pakes and McGuire generate stochasticity from the random realization of investment, rather than from random investment costs.²¹ Therefore, in Pakes and McGuire, quality investments are deterministic conditional on observable state variables while investment outcomes are not. Computation of Pakes and McGuire involves an exact solution for the level of investment at each state (Pakes et al., 1992). Researchers have extended the model by developing a closed form for the density of investment (Borkovsky et al., 2017). Their methods would not directly apply to our case, because the fact that each investment level results in a different future state in our model leads to difficulties in specifying a closed form density and to some investment levels never being chosen in a discrete approximation.

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.²²

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 shipbuilding in a dynamic oligopoly model. Each period

²¹To ensure existence of pure strategy equilibrium, their model satisfies the unique investment choice (UIC) admissibility criterion, as defined by Doraszelski and Satterthwaite (2010) and extended by Escobar (2013) to incorporate multidimensional investment decisions.

²²A 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.

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 she writes $VE(s_0^t)$. Production cost has a deterministic component that is a function of the amount produced and the shipyard’s backlog, $\bar{c}_i(a_i^t, s_i^t, s_0^t)$.²³ Production cost also has a stochastic normally-distributed component with standard deviation σ^t . The period payoff is thus given by:

$$\underbrace{VE(s_0^t) \times a_i^t}_{\bar{\pi}_i(a_i^t, s^t)} - \bar{c}_i(a_i^t, s_i^t, s_0^t) - \underbrace{a_i^t \sigma_i^t}_{\tilde{c}_i(a_i^t, s^t)} \times \varepsilon_i^t, \quad (10)$$

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 $\tilde{c}_i(a_i^t, s^t) \times \varepsilon_i^t$. Kalouptsidi (2018, p. 1123) shows that convexity of the cost function $\bar{c}_i(a_i^t, s_i^t, s_0^t)$ 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.²⁴ Our simulation evidence in Section 5 below shows that this the assumption does not hold in some real-world empirical settings. Our results complement the Kalouptsidi results by showing how to compute optimal policies for general payoff functions which do not necessarily induce discrete concave choice-specific values.

²³Kalouptsidi (2018) refers to this term as c , but we use our notation of \bar{c}_i instead. Similarly, we adjust the notation of other papers below.

²⁴Koh (2022) makes a similar assumption.

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, $\bar{\pi}_i(h^t, s_i^t)$.

The period payoff is given by:

$$\bar{\pi}_i(h^t, s_i^t) - \underbrace{a_i^t \times p^t}_{\bar{c}_i(a_i^t, s_i^t)} - \underbrace{a_i^t}_{\bar{c}_i(a_i^t, s_i^t)} \times \varepsilon_i^t. \quad (11)$$

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. Caoui estimates the choice probabilities with simulation methods.

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. It then earns profits by procuring and selling electricity over the period, earning a regulated rate-of-return that declines in its variable costs. 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, \tilde{c}_1 is proportional to the capacity added or retired, implying that the stochastic component of cost is linear in the amount invested. For either generation source, we can write the action a_1^t as the next period's capacity level of that source.

Our model considers a single investment decision, but in Gowrisankaran et al. (2024), 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 period payoffs for adding gas capacity (or retiring coal capacity) can then be written as:

$$\bar{\pi}_1(s_1^t, s_0^t) - \underbrace{\delta_0 \mathbb{1}\{a_1^t \neq s_1^t\} + (a_1^t - s_1^t) (\delta_1 + (a_1^t - s_1^t)\delta_2)}_{\bar{c}_1(a_1^t, s^t)} - \underbrace{(a_1^t - s_1^t)\sigma}_{\tilde{c}_1(a_1^t, s^t)} \times \varepsilon_1^t, \quad (12)$$

where $\bar{\pi}_1(s_1^t, s_0^t)$ are the profits from operations model for the CCNG decision, or 0 for the coal decision. 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. Gowrisankaran et al. found that, for some states, 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, $\bar{c}_i(a_i, s_i)$. The stochastic part of the investment cost, $\tilde{c}_i(a_i, s_i)$, is proportional to investment, $a_i^t - s_i^t$.²⁵ 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:

$$\begin{aligned}
\bar{\pi}_i(s^t) &= \mathbb{1}\{a_i^t > s_i^t\} (\delta_1 + \delta_2(a_i^t - s_i^t) + \delta_3(a_i^t - s_i^t)^2 + \mathbb{1}\{s_i^t = 0\}\chi) \\
&- \mathbb{1}\{a_i^t < s_i^t\} (\delta_4 + \delta_5(a_i^t - s_i^t) + \delta_6(a_i^t - s_i^t)^2 - \mathbb{1}\{a_i^t = 0\}\phi) \\
&- [\mathbb{1}\{a_i^t < s_i^t\}(a_i^t - s_i^t)\sigma_1 + \mathbb{1}\{a_i^t > s_i^t\}(a_i^t - s_i^t)\sigma_2] \times \varepsilon_i^t.
\end{aligned} \tag{13}$$

In equation (13), the first two lines make up the deterministic component of period payoffs, $\pi_i(a_i^t, s^t)$. 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 $\tilde{c}_i(a_i^t, s^t) \times \varepsilon_i^t$. 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 Fowlie 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 is deterministic. We next turn to simulation of

²⁵Our working paper on rural hospitals, Gowrisankaran et al. (2010), used this framework.

models like Gowrisankaran et al. (2024) to further understand the impact of the unobservable distribution on outcomes.

5 Simulation evidence

We provide simulation evidence on the economic implications of alternative modeling choices and the performance of our model as the number of choices increases. To do this, we use data and estimates from the application by Gowrisankaran et al. (2024) described in Section 4.4 above, which we refer to as GLR henceforth. We are interested in understanding the impact of two alternative modeling choices for the marginal cost of investment.

First, we consider the model used by GLR and analyzed here, which has a linear shock to the marginal cost of investment, as in equation 1. While GLR consider both coal retirement and CCNG investment, we focus on new CCNG capacity. Our main results use the model parameters from GLR, Table 6. We use the same grid for potential actions a_1 , with one departure. Whereas GLR allow for 10 CCNG capacity investment levels that range from 0 to 3000 MWs (and coal capacity retirement levels that generally range from 0 to 5000) but are more concentrated towards 0, we allow for evenly spaced retirement/investment levels over these same ranges. This then allows us to examine the impact of varying the number of investment levels, keeping the spacing between them the same. We simulate each experiment 10,000 times and use the same data and 10-period decision horizon as in GLR.

Second, we analyze a model where an *i.i.d.* shock is added to each capacity choice. This model is also widely used by papers analyzing investment choices (e.g., Chatterjee et al., 2023). In GLR, the residual is normal and is scaled by estimated parameters σ^f , for fuel/technology type $f \in \{CCNG, coal\}$. Here, we use *i.i.d.* type 1 extreme value residuals, one for each investment choice. An important question is the scale of these residuals. For CCNG, we scale the logit residuals by a factor of $\sigma^{f,MNL} = 333.3 \times \sqrt{3}/\pi \times \sigma^f$ (using 555.6 for coal). Starting from the end, the third term is the same as in GLR and the second term scales divides by the standard deviation of the type 1 extreme value distribution. Turning to the first term, 333.3, while there is no way to exactly replicate the variance in the linear

marginal cost shock with *i.i.d.* shocks, this value scales each *i.i.d.* shock by the standard deviation of the difference between subsequent grid points.

Table 1: Base simulations

# inv. choices	Mean investment (MW)		Std. dev. investment (MW)		Mean value (millions of 2006 \$)			Mean # zero prob. choices
	GSD	MNL	GSD	MNL	GSD	MNL	AR	GSD
5	262.3	257.7	711.5	598.7	3,258	3,284	2,344	0.1
10	263.5	319.8	708.3	574.9	3,306	3,578	2,204	0.4
15	263.5	342.4	708.9	563.7	3,318	3,776	2,143	0.6
20	264.1	355.0	708.5	557.5	3,318	3,926	2,104	1.1
25	264.0	363.0	708.1	553.9	3,320	4,051	2,082	1.5
30	264.9	368.1	708.6	551.2	3,322	4,154	2,068	1.9
35	264.6	371.3	708.4	548.3	3,323	4,246	2,057	2.5
40	264.5	374.3	707.9	547.0	3,323	4,325	2,049	2.7
50	264.9	378.4	708.6	544.9	3,325	4,462	2,034	3.7
60	264.9	381.2	708.4	543.7	3,325	4,576	2,027	4.6
70	264.8	383.1	708.3	541.8	3,326	4,673	2,021	5.4
80	265.4	384.4	708.4	541.2	3,326	4,759	2,017	6.3

Note: GSD (the authors' initials) refers to the model considered in this paper, which has shocks to the marginal cost of investment. MNL refers to the model with *i.i.d.* multinomial logit cost shocks. AR refers to the MNL model with the Akerberg and Rysman (2006) correction. Computations from GLR sample and parameters with 10,000 simulation draws.

Table 1 presents base simulation results for both models. Considering first the mean and standard deviation of investment in columns 2 and 4, we find that our model converges to very similar values for 10 or more investment choices, with a mean investment level of about 265 MWs per period and a standard deviation of about 708 MWs. Recall that GLR chooses discrete investment levels as an approximation to a continuous choice of investment. The fact that the discrete approximation exhibits convergence provides reassurance that the discrete choice model with linear shocks to the marginal cost of investment and discrete investment

levels approximates a continuous choice model. With 10 capacity choices, our exact and fast algorithm requires fewer than 20 comparisons while a naïve algorithm would require 45 comparisons, with the computational cost of the naïve algorithm growing quadratically in the number of choices, but only linearly for our algorithm.

Turning to the model with multinomial logit shocks in columns 3 and 5, we find that its investment levels do not converge over the range of investment choices that we tried. While the multinomial model with 5 choices provides reasonably similar investment statistics to our model with 5 choices, the results between the two models diverge as the number of choices increases. In particular, the mean investment level continues to increase while the standard deviation of investment decreases (while our model exhibits stable patterns). With many options, the multinomial model will more resemble picking a random option in any region, because there will more likely be an option with a high *i.i.d.* draw near any investment level. However, the random draws will be weighted so that the firm will want to pick a draw that is close to the one with the highest choice-specific value function.

We next consider Table 1 columns 6 and 7, which show the mean *ex-ante* values, V_1 , that result from optimizing decision-making across the two models. The model with linear shocks to the marginal cost of investment finds mean values that converge to about \$3.325 billion as the number of choices increases. As with investment decisions, we find very little variation with more than 10 investment choices. We know theoretically that the values for the multinomial model will grow at a log rate as the number of choices increases, and hence that it will asymptote to infinity. Our simulation results for the multinomial model verify that the values continue to grow as the number of choices increases.²⁶

The lack of convergence for the multinomial model shows that it cannot well approximate a continuous choice model and may lead to biased predictions of long-run market structure. It may also lead to inconsistent parameter estimates when used in a structural estimation process. For instance, the fact that it overpredicts investment may lead to an upwardly

²⁶Akerberg and Rysman (2005) suggest a correction to the static multinomial model, that applied to our context implies subtracting $\sigma^{f,MNL} \times \log(K)$ from V_1 . Although the structural interpretation of their correction is not applicable to settings with capacity choices, Table 1 column 8 nonetheless computes a variant of our model with the correction. As with the base multinomial model, the variant does not generate stable values in our setting. It has values that decline as the number of choices increases.

biased estimate of investment fixed costs. The high predicted *ex-ante* values generated by the logit draws may also lead to an overestimate of the sunk entry costs.

Finally, we turn to the question of whether all investment levels are chosen with positive probability in our model. Table 1 column 9 shows the mean number of levels that are never chosen along the 10-period equilibrium path across sample utilities for our model.²⁷ In the base model with 10 investment levels, on average, 0.43 levels are never chosen. This suggests that assuming that all choices are on the discrete concave hull may not be a good approximation in empirical settings where period profits, and hence choice-specific value functions, vary in complex ways.²⁸ We also find that the fraction of levels that are never chosen increases more than proportionally with the number of choices. To see how this might happen, consider the example in Figure 2 which has 3 actions, one of which is never chosen. If we went from 3 to 5 actions, adding 1.5 and 2.5 as potential levels with choice-specific value functions that are a linear combination of their neighbors, these choices would both be dominated by the extremal actions $k = 1$ and $k = 5$. Thus, in this case, the number of dominated actions would go from 33.3% to 60%. Finally, our results suggest that one may want to use GMM rather than maximum likelihood to estimate this class of models because the presence of zero probability events may lead to zero likelihoods.

Some of the investment results in Table 1 may reflect the exact estimated parameters from GLR. In particular, GLR estimated small and slightly negative fixed costs of investment. This stands in contrast to manufacturing industries, which typically have lumpy investments and substantial shares at zero (Cooper and Haltiwanger, 2006), implying large and positive fixed costs. To understand the implications of fixed costs, Table 2 provides simulations from the same model, but setting the fixed cost parameters δ_1^{CCNG} and δ_1^{COAL} so that the fixed costs of these investments are equal to the linear mean marginal costs of a 300 MW generator, roughly the median size of a new generator.

As with the base specification, we find that our model is stable with a sufficient number of investment choices. However, with this specification, the mean and standard deviation of

²⁷In the model with multinomial shocks, all options are chosen with positive probability by assumption.

²⁸In GLR, the period profits $\bar{\pi}$ stem from a monopoly utility's optimization of a dynamic operations model when faced with multiple and overlapping regulatory incentives.

Table 2: Simulations with significant fixed costs

# inv. choices	Mean investment (MW)		Std. dev. investment (MW)		Mean value (millions of 2006 \$)			Mean # zero prob. choices
	GSD	MNL	GSD	MNL	GSD	MNL	AR	GSD
5	179.6	149.5	651.6	547.1	3,095	3,079	2,116	2.2
10	197.6	190.1	684.8	558.3	2,698	2,778	1,412	5.8
15	200.5	216.0	689.9	564.9	2,726	2,881	1,273	9.8
20	205.3	237.5	697.8	571.3	2,762	2,977	1,207	13.6
25	206.0	252.1	699.5	573.6	2,778	3,051	1,146	17.4
30	205.9	263.2	698.5	574.2	2,783	3,107	1,089	21.4
35	206.2	271.9	699.7	574.2	2,798	3,169	1,056	25.1
40	206.6	279.6	700.6	573.7	2,808	3,217	1,027	29.3
50	206.3	291.9	699.8	572.8	2,819	3,306	981	36.8
60	206.6	301.1	700.3	571.4	2,825	3,380	939	44.9
70	206.5	308.6	700.2	570.1	2,835	3,447	911	52.7
80	206.8	314.7	700.8	568.5	2,837	3,502	885	60.3

Note: GSD (the authors' initials) refers to the model considered in this paper, which has shocks to the marginal cost of investment. MNL refers to the model with *i.i.d.* multinomial logit cost shocks. AR refers to the MNL model with the Akerberg and Rysman (2006) correction. Computations use GLR sample and parameters with 10,000 simulation draws but modify FC to be equal to the linear MC of a 300MW plant.

investment converge with 20 or more choices, while 10 choices was sufficient with the base parameters. Given the higher fixed costs of CCNG investment, it is not surprising that our model predicts a lower level of investment than with the base GLR parameters, with mean investment per period dropping from 265 MWs to 206 MWs. The standard deviation of investment goes down, but only slightly, from 708 MWs to 700MWs.

The number of investment levels that are never chosen increases dramatically to the majority of choices for all simulations with more than 5 investment choices. This is because it is no longer profitable for the firm to add a small amount of capacity for *any* ε . If ε is sufficiently high, it will invest nothing, and when ε drops enough to overcome the fixed cost of investment, the firm will add a substantial amount of capacity, skipping the small investment levels.

Considering the model with multinomial shocks, we find again that the mean and standard deviation of investment do not converge as the number of choices increases. Moreover, the mean investment increases at a faster rate than with the base GLR parameters, going from about 150 to 315 MWs versus 260 to 380 MWs with the base parameters. This occurs because there is only one draw where the firm will not have to pay the fixed cost—no investment—and so the proportion of draws where the firm pays the fixed cost increases with the number of choices. Thus, it will be more likely to invest a positive amount with more draws. This feature is driven by the *i.i.d.* nature of the residual structure and does not occur with our model.

The increase in mean investment with more choices with the multinomial model also leads to an increase in the standard deviation of investment with more choices. This is the opposite of the multinomial model estimated with the base GLR parameters. Thus, it shows that the biases from using a multinomial model when the data are generated from a model with linear shocks to the marginal cost of investment will be difficult to sign. Overall, this table confirms our general point that the multinomial model lacks stability as the number of choices increases, while our model converges with a sufficient number of grid points.

6 Conclusion

This paper develops new methods that are useful in computing and estimating a class of dynamic oligopoly models. We consider models where firms can invest to build or retire capacity or other attributes, and where the potential investment actions have many discrete values that can approximate a continuous distribution. The desired investment levels are typically limited by convex adjustment costs and capital specificity. This class of models is quite general and has been used repeatedly in the recent literature, in settings ranging from electricity utilities building new generators to shipbuilders adding to their order queue.

We argue that private information shocks are useful in this setting, both for ensuring existence of pure strategy equilibrium when there are multiple firms and for generating variation in predicted choices as a basis for empirical work. We consider a functional form for the unobservable that we believe is useful in these models: private information shocks to the marginal cost of investment. This functional form has been used by an important subset of papers within this class of models.

We characterize the optimal investment policy with this type of shock. Specifically, Proposition 1 shows that investment options are chosen with positive probability for some value of the private information shock if and only if they are in the discrete analog of the concave hull of the choice-specific value function relative to other actions. Proposition 2 uses these results to develop a computationally quick method to calculate the probability of choosing each action. Our method requires $2K - 3$ or fewer comparisons per state and firm, where K is the number of choices—i.e. potential investment levels—while a naïve algorithm would require $K(K - 1)/2$ comparisons, and hence is slower for $K > 3$.

Previous papers have solved for equilibria either by simulating the unobservables or by assuming functional forms for which all choices are in the discrete concave hull. We show with simple examples that approaches that simulate the unobservables may lead to difficulties in ensuring the existence of equilibrium. We also simulate data from one empirical study that uses our approach, Gowrisankaran et al. (2024), finding that a significant portion of choices are not in the discrete concave hull. Thus, we believe that in many complex empirical

problems, the choice-specific value functions may not lie in the discrete concave hull.

Other papers in the literature have assumed multinomial *i.i.d.* unobservables for each discrete action. We provide simulation evidence that, with these models, the investment choices and ex-ante values are not stable in the number of choices. In contrast, the simulation evidence shows that our model converges to very similar ex-ante values and investment choices with sufficient grid points.

Overall, we believe that models with a continuous or near-continuous investment choice with unobservable shocks to the marginal cost of investment are important in a variety of economic settings in industrial organization, energy economics, and macroeconomics. Our method provides a useful extension of the empirical toolbox that provides a quick way of computing these models and estimating these models with GMM estimation, potentially combined with conditional choice probability methods. We provide Python code—at <https://github.com/patohdzs/gsd-capacity-investment>—that researchers can use to easily implement our algorithm.

References

- Akerberg, D. and Rysman, M. (2005). Unobserved product differentiation in discrete-choice models: Estimating price elasticities and welfare effects. *Rand Journal of Economics*, 36(3):738–761.
- Aguirregabiria, V., Collard-Wexler, A., and Ryan, S. P. (2021). Dynamic games in empirical industrial organization. In *Handbook of industrial organization*, volume 4, pages 225–343. Elsevier.
- Aradillas-López, A. and Gandhi, A. (2016). Estimation of games with ordered actions: An application to chain-store entry. *Quantitative Economics*, 7(3):727–780.
- Bajari, P., Benkard, C. L., and Levin, J. (2007). Estimating dynamic models of imperfect competition. *Econometrica*, 75(5):1131–71.

- 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. (2010a). Learning-by-doing, organizational forgetting, and industry dynamics. *Econometrica*, 78(2):453–508.
- Besanko, D., Doraszelski, U., Lu, L. X., and Satterthwaite, M. (2010b). Lumpy capacity investment and disinvestment dynamics. *Operations Research*, 58(4-part-2):1178–1193.
- Borkovsky, R. N., Goldfarb, A., Haviv, A. M., and Moorthy, S. (2017). Measuring and understanding brand value in a dynamic model of brand management. *Marketing Science*, 36(4):471–499.
- Caoui, E. H. (2023). Estimating the costs of standardization: Evidence from the movie industry. *The Review of Economic Studies*, 90(2):597–633.
- Carroll, C. D. (2006). The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics letters*, 91(3):312–320.
- Chatterjee, S., Corbae, D., Dempsey, K., and Ríos-Rull, J.-V. (2023). A quantitative theory of the credit score. *Econometrica*, 91(5):1803–1840.
- 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.
- Escobar, J. F. (2013). Equilibrium analysis of dynamic models of imperfect competition. *International Journal of Industrial Organization*, 31(1):92–101.
- 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., Lucarelli, C., Schmidt-Dengler, P., and Town, R. (2010). Government policy and the dynamics of market structure: Evidence from critical access hospitals. Working Paper.
- Gowrisankaran, G. and Town, R. J. (1997). Dynamic equilibrium in the hospital industry. *Journal of Economics and Management Strategy*, 6:45–74.
- Granja, J. and Forte, G. (2024). Recursivity and the estimation of dynamic games with continuous controls.
- Hotz, V. J. and Miller, R. (1993). Conditional choice probabilities and the estimation of dynamic models. *Review of Economic Studies*, 60:497–529.

- Iskhakov, F., Jørgensen, T. H., Rust, J., and Schjerning, B. (2017). The endogenous grid method for discrete-continuous dynamic choice models with (or without) taste shocks. *Quantitative Economics*, 8(2):317–365.
- 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.
- Koh, P. S. (2022). Estimating discrete games of complete information: Bringing logit back in the game. *arXiv preprint arXiv:2205.05002*.
- 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. (1997). Using randomization to break the curse of dimensionality. *Econometrica: Journal of the Econometric Society*, pages 487–516.

- 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.
- Snashall-Woodhams, N. (2024). Targeting solar subsidies. Technical report, Working Paper.
- Srisuma, S. (2013). Minimum distance estimators for dynamic games. *Quantitative Economics*, 4(3):549–583.