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 October 2026

We thank Patricio Hernandez, Filippo Paternollo, and Hara Schwartz for excellent research assistance. We thank Hassan Afrouzi, El Hadi Caoui, Allan Collard-Wexler, Michael Dinerstein, Jean-Fran ois Houde, Mitsuru Igami, Fedor Iskhakov, Paul Koh, Ashley Langer, anonymous referees, and seminar and conference participants at the University of Arizona, Cornell University, CREST-INSEE, Econometric Society World Congress, Universita 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, Schmidt-Dengler from the Austrian Science Foundation (FWF, project FG 5 TP 5). 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.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 October 2026
JEL No. C57, C63

ABSTRACT

We consider widely used dynamic models where every period, each agent receives a private information shock to investment marginal cost and chooses among many ordered investment levels. We provide a computationally fast method to calculate optimal investment probabilities given value functions. Our method finds shock cutoffs between investment choices, accounting for the possibility that some levels are dominated. Other methods to compute these models may result in non-existence of pure strategy equilibrium or only work under restrictive assumptions. Multinomial logit models with choice-specific shocks can lead to counterintuitive predictions. Our method is useful to researchers estimating dynamic models of investment.

Gautam Gowrisankaran
Columbia University
Department of Economics
and CEPR
and also NBER
gautamg2@gmail.com

Philipp Schmidt-Dengler
University of Vienna
and CEPR
and CES-Ifo
and WIFO
philipp.schmidt-dengler@univie.ac.at

1 Introduction

Many economic settings involve agents making decisions that are hard to reverse. Policies often interact with these decisions, for instance by subsidizing particular actions. Irreversible actions are inherently *dynamic* in that they require the agent to understand the expected future value of committing to the decision relative to waiting to acquire additional information. To analyze these settings, over the past four decades, researchers have developed a computational and empirical literature analyzing dynamic models, often with many endogenously heterogeneous agents interacting (Ericson and Pakes, 1995).

Often in these models, agents differentiate themselves through investment decisions with a natural order that affects their future state, which we can generically think of as capacity. These investment decisions frequently have many choices which might also approximate a continuous set of choices. Empirical papers across industrial organization, macroeconomics, and energy economics have studied contexts 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 (Kalouptsidi, 2018), individual asset level changes (Chatterjee et al., 2023), digital movie screen adoption (Caoui, 2023), electricity capacity investments (Gowrisankaran et al., 2026), and rooftop solar installations (Snashall-Woodhams, 2024). An important issue in these models is how to specify the unobservable components of the values across actions, which we can think of as shocks to the cost of investment.

This paper provides three insights as to the modeling of this common dynamic problem. First, we argue that having one *i.i.d.* private information shock to the marginal cost of investment often results in a more attractive specification than no shock, a shock to each option, or a shock to the fixed cost of investment. Second, we show that many important economic forces, such as investment fixed costs, will lead this class of models to have options that are chosen with zero probability in equilibrium and that this complicates estimation and computation. Third, we offer a computational algorithm to compute the choice probability of each option for this class of models.

Our algorithm deals with options that are never chosen and is fast, scaling linearly in the number of choices. The output of this algorithm is the set of investment choices that are chosen with positive probability, the cutoff values of the unobservable that make the agent indifferent between two adjacent choices, and the probability each option is chosen. We start by showing that an investment level will be chosen with positive probability if and only if it lies in the discrete analog of a convex hull of the choice-specific value function. We then develop a method to find the set of choices that are in this discrete convex hull. Our method iteratively compares an investment choice against its neighbor, marking a choice as having zero probability if the cutoff value is higher than the previous cutoff, when it should be lower. In this instance, the algorithm works backward, potentially eliminating earlier choices based on the same criterion. We provide simple Python code to implement the algorithm.[1]

We believe that our model and algorithm are applicable to many dynamic estimation and computation papers that consider investment, broadly defined. Focusing first on our *model*—i.e., a private information shock to the marginal cost of investment—many investment papers use private information shocks as they generate heterogeneity in decisions for a given state, which is important as a basis for empirical work (Rust, 1987). The shocks also make *ex-ante* reaction functions continuous, which is used to ensure existence of pure strategy Markov Perfect Equilibrium and to compute dynamic game equilibria with multiple agents (Gowrisankaran, 1995; Doraszelski and Satterthwaite, 2010).

Perhaps the most common shock specification is an *i.i.d.* type 1 extreme value shock to each option, which results in a multinomial logit (MNL) model.[2] However, the MNL model does not fit the economics of many problems in that firms typically do not receive independent shocks for two similar capacity choices. This is particularly problematic because the number of investment levels needs to be large to approximate continuous investment and is often

[1]See `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.

[2]For instance, in their empirical study of the nuclear power plant industry, Rust and Rothwell (1995) model the capacity decision of nuclear power plants with an *i.i.d.* logit shock to each choice of capacity utilization. More recent papers that use logit shocks include Collard-Wexler (2013), Chatterjee et al. (2023), Snashall-Woodhams (2024), and Wang (2026).

chosen arbitrarily. As we illustrate with empirical simulation evidence in Section 5, the predictions of the MNL model may not be robust to the number of levels.

Turning to our *algorithm*, it offers advantages over two other common algorithms with the same shock specification. First, 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). The issue with this approach is that, with a finite number of simulation draws, action probability best responses will generically not be continuous in the parameters, which may lead to non-existence of equilibrium (as we illustrate in Section 4.1). Even in single agent models, where existence of a pure strategy equilibrium is generally not an issue, the probabilities recovered by a simulation approach are step functions, making it difficult to estimate the parameters with a non-linear search algorithm. Second, one can leverage concavity assumptions on the period payoff functions to ensure that every option is chosen with positive probability (Kalouptsidi, 2018). This method is useful in some settings but these assumptions are typically inconsistent with fixed costs of investment, quality ladder investments (Pakes and McGuire, 1994), and dynamic game interactions, all of which are important in many other economic contexts.

Accordingly, we recommend that scholars use our algorithm in the following two parts of an empirical dynamic paper. First, we suggest using it to compute counterfactual equilibria. Since our algorithm considers the calculation of best response probabilities, researchers may want to nest it within an iterated best response (as used by Pakes and McGuire, 1994) or other method to solve for a new equilibrium. Second, we also suggest using our algorithm for estimation, because it recovers choice probabilities that are continuous in parameters. As having shocks to the marginal cost of investment implies that some events may be chosen with zero probability, researchers should use a generalized method of moment (GMM) estimation method rather than maximum likelihood. In a full solution estimation approach, researchers can compute the equilibrium for each parameter vector to evaluate the GMM criterion. Researchers can also use a conditional choice probabilities (CCP) method (of the type suggested by Pakes et al., 2007, for instance), obtaining choice-specific value functions

for candidate parameters, solving for the probability of each action with our method, and then embedding these probabilities in the GMM criterion.[3]

The remainder of the paper is structured as follows. Section 2 describes our modeling framework, including discussing existence and uniqueness of equilibrium. Section 3 discusses the economic implications of alternative specifications for investment costs. It covers both the specification discussed in this section and others, and explains the modeling assumptions used by different papers in the literature. Section 4 specifies our algorithm and the accompanying theoretical results that we use to justify it. Section 5 provides simulation evidence comparing outcomes for an empirical paper that uses shocks to the marginal cost of investment to a version of that paper with scaled *i.i.d.* MNL shocks on each option. Section 6 concludes.

2 Model and Equilibrium

2.1 Model

We consider investment in a dynamic oligopoly framework with discrete time, $t = 1, 2, \ldots, \infty$. 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 $\mathcal{N} = \{1, ..., N\}$, 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 levels.

We note two general points regarding our model. First, although our model formally considers the case of a fixed number, K, of levels, our approach may also be useful as an approximation to the equilibrium as the number of discrete levels increases.[4] 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

[3] Another branch of this literature has emphasized one-step ahead estimators that obtain relative values by inverting CCPs (Hotz and Miller, 1993; Arcidiacono and Miller, 2011). In our case, the presence of options chosen with zero probability means that this mapping is no longer unique, implying that these estimators cannot be applied without modification.

[4] An important caveat is that we did not evaluate the limiting equilibrium with continuous choices. It is well known that this type of limit may behave irregularly or produce new equilibria and thus this exercise is beyond the scope of this paper.

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 $s^t \in S \subseteq \Re^L$. The state includes the capacity level of each firm and potentially other firm-specific characteristics. It also includes common characteristics such as aggregate demand or aggregate productivity. We 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 (i.e., the firm's capacity) and s_0^t denote the portion that is unaffected by the actions of any firm (e.g., market demand).

Privately observed cost shocks. At time t, each firm i privately observes a real-valued cost shock $\varepsilon_i^t \in \Re$. 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 $\Re$. 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 $\mathbb{E}\left[\varepsilon_i^t | \varepsilon_i^t \geqslant \varepsilon\right] < \infty$ 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 $A_i = (\alpha^1, \ldots, \alpha^K)$; i.e. A_i consists of K unique real numbers.[7] Using $o(\cdot)$ to denote the cardinality of sets, $o(A_i) = K$. We impose an increasing order on A_i, so $\alpha^1 < \alpha^2 < \ldots < \alpha^K$, and let an action profile a^t denote the vector of joint actions in period t, $a^t = (a_1^t, \ldots, a_N^t) \in A = \times_{i=1}^N A_i$. The cardinality of the action space A is given by K^N.

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

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

[7] For ease of notation, we restrict $\alpha^k, \forall k$ and K to be identical across agents but our approach can accommodate having these elements be player-specific.

State transitions. We describe the state transition matrix with a probability density function $g : A \times S \times S \longrightarrow [0,1]$ where a typical element $g\left(s^{t+1}|a^t, s^t\right)$ is the probability that state s^{t+1} is reached when the current action profile and state are given by (a^t, s^t). We require $\sum_{s' \in S} g\left(s'|a, s\right) = 1$ 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. Unlike Pakes and McGuire (1994), our framework allows for firm i's observable state variable s_i^t to evolve deterministically (but with *ex-ante* stochastic costs), as might occur in settings such as capacity accumulation. 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:

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

where $\tilde{c}_i(a_i^t, s^t)$ is the component of investment cost that multiplies the unobservable shock. 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.

Empirical applications we discuss all use the special case of shocks to the marginal cost, in which case $\tilde{c}_i(a_i^t, s^t) \equiv a_i^t - s_i^t$.

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. We can separate $\pi_i(a_i, s)$ into these two components: $\pi_i(a_i, s) \equiv \overline{\pi}_i(a_i, s) - \overline{c}_i(a_i, s)$. A natural specification for $\overline{c}_i(a_i, s)$ would be

[8]For 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).

a quadratic total cost function with three parameters, on fixed costs, quantity, and quantity squared. We assume that both π_i and $\tilde{c}_i$ are 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:

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

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

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

obtain:

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

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 exposit 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

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

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:

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

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

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

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

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

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 an MPE, beliefs ω must correspond to choice probabilities p so that (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. 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.

While an MPE exists, there is no guarantee of uniqueness. However, the methods suggested in this paper can also be applied when alternative timing assumptions are made in order to guarantee uniqueness. For instance Igami and Uetake (2020) and Elliott (2026) assume a finite horizon and sequential moves to ensure uniqueness in their models.

3 Economics of Alternative Cost Specifications

This section first discusses the economics of models with unobservable shock on investment marginal costs. It then turns to other models. For both sets of models, we explain the relation of our model to papers in the literature.

3.1 Models with an Unobservable on Marginal Cost of Investment

As developed in Section 2.1, we propose a model with private information shocks to the marginal cost of investment. This specification has several economic advantages over alternatives. First, having a private information shock ensures existence of pure strategy equilibria. Second, it generates stochasticity in investment outcomes, as the investment is decreasing in the realization of the shock. Third, a shock to the marginal cost of investment has a natural interpretation as an unobserved component of the cost of installing a unit of capacity, which might be caused by random variation in the price of an input, for instance. Finally, our approach allows for lumpy investment that is not restricted by the current installed capacity, i.e. firms can make infrequent but large jumps rather than being restricted to changing the attribute by small amounts each period in order to adjust. This might be important in, e.g., manufacturing industries, where investment can be lumpy (Cooper and Haltiwanger, 2006).

For some of these economic reasons, a number of papers use the approach that we highlighted in this paper. Specifically, these papers model capital stock, broadly defined, as the main state variable, s_{it}. They allow each firm to choose any level of capital stock that it wants with investment. Investment costs include a term that is linear in the amount invested, and sometimes fixed and quadratic terms. In addition, each firm faces a private information *i.i.d.* shock to the marginal cost of investment.[10]

Papers that fit within this framework include the following works. Kalouptsidi (2018) studies the impact of subsidies in the market for shipbuilding. Her time-varying state is the number of ships in the queue. Barwick et al. (2025) also study shipbuilding but focus on

[10]Mathematically, our decision problem is also equivalent to the classic Bresnahan (1987) vertical demand system, where the probability a firm chooses a given investment level in our model is equivalent to the market share of a particular automobile in their case. In both contexts, zero share options may exist.

the shipyard's capital stock as the state variable. Caoui (2023) studies adoption of digital movie screens. His state variable is the number of digital screens owned by any company. Investment corresponds to converting more screens to a digital format. Gowrisankaran et al. (2026) study regulated monopoly electric utilities faced with an energy transition. Their state variables include the MWs of coal and combined-cycle natural gas (CCNG) capacity. Firms can retire coal or invest in CCNG each period. Finally, Yun (2025) studies whaling firms in Nantucket, focusing on ship capacity as the state variable. These papers have used a variety of techniques for estimation, including specifying economic fundamentals that result in having no dominated options (Kalouptsidi, 2018), simulating choice probabilities (Caoui, 2023), approximating value functions with polynomials (Barwick et al., 2025), and applying the methods of this paper (Gowrisankaran et al., 2026; Yun, 2025).[11]

As we previewed in Section 1, having a choice-specific value function that is not concave when applied to the discrete points in A results in some investment levels being chosen with zero probability. Alternately put, this means that no realization of ε_i makes these investment levels the optimal choice.

But, when should we worry about non-concavity economically? In some settings, such as a capacity constrained monopolist with linear demand and constant marginal costs, profits are concave in capacity. Together with convex investment costs, this makes the choice-specific value functions concave in investment, which then implies that all options are chosen with positive probability (see Kalouptsidi, 2018, Lemma 1). Yet, there are at least three important economic fundamentals that can cause this framework to break down.

First, capacity adjustments often involve fixed costs. For instance, the fact that manufacturing industries generally exhibit lumpy investment (Cooper and Haltiwanger, 2006) suggests substantial fixed costs. Such fixed costs will cause non-convexities in the investment costs and hence likely result in a choice-specific value function that is not concave.

Second, period profits may themselves be convex in the choice variable. A prominent example is the flow profit function used by Pakes and McGuire (1994). Profits here arise

[11] Relatedly, an econometrics literature extends the estimation of dynamic oligopoly models by leveraging assumptions that agents' choice belongs to the discrete convex hull (Srisuma, 2013; Aradillas-López and Gandhi, 2016; Koh, 2022).

from a logit demand system, where consumer utility is linear in quality and firms can choose quality. The monopolist's profits are increasing in quality at an increasing rate. Hence the profit function is convex in quality.[12] This will also likely result in a choice-specific value function that is not concave.

A final source of non-concavity is the strategic interactions in dynamic oligopoly. Standard restrictions on the period payoff function that yield concavity of the value function may no longer be sufficient in this case because a firm's actions may affect its rivals' future actions and in turn its own payoffs. As an example, a moderate investment level might provoke a rival to invest more to compete, while a high investment level might cause this rival to give up and exit. This might mean that the firm would then never choose a moderate investment level. Alternately put, this could generate a choice-specific value function that is not concave, even with convex investment costs. Even in instances where parameter estimates show that investment returns are concave, counterfactual policies such as subsidies for investments over a threshold would introduce non-concavities.

3.2 Models with Other Investment Cost Specifications

While we proposed a model with shocks to the marginal cost of investment in Section 2.1, other sources of stochasticity in investment outcomes have been used in the ordered choice context that we consider. One approach has been to completely omit a private information unobservable. The macroeconomics literature has employed this method since in many cases it considers single agent or competitive settings and is less concerned with fitting data where agents take different actions at the same observable state. The endogenous gridpoints method (Carroll, 2006) shows how to optimize computation of this type of model and this method has been extended to discrete-continuous settings (Iskhakov et al., 2017).

A common other approach to an unobservable shock to the marginal cost of investment is an $i.i.d.$ shock to each action a_{ik}, which we can call ε_{ik}. As noted in Section 1, the most common specification for ε_{ik} has been type 1 extreme value, resulting in the MNL model.[13]

[12] Derivation available from the authors upon request.

[13] Mermelstein et al. (2020), which studies mergers computationally, specifies an alternative, which is a

Having *i.i.d.* action-specific cost shocks results in continuous reaction functions which will ensure existence of equilibrium. However, the results of this specification may not be robust to the number of investment choices, K, a point we substantiate in Section 5. This model also has the undesirable property that a 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. As an alternative to *i.i.d.* shocks on each option, some authors have used shocks to the fixed cost of investment (Khan and Thomas, 2008; Ryan, 2012; Fowlie et al., 2016). In this case, the ε_i term is paid only if the firm alters its capacity. A limitation of this specification is that it does not ensure equilibrium existence with finite investment levels because the chosen investment level can jump in other firms' actions. It also does not generate heterogeneity in outcomes conditional on a non-zero investment level. Generating such heterogeneity is likely to be important in contexts where the size of investment is a meaningful strategic decision.

A final approach has been to limit the set of states that can be reached in one period. The canonical Pakes and McGuire (1994) model only allows a firm's product quality to vary by 1 between consecutive periods. Other papers that use this approach in empirical work include Igami and Uetake (2020), Caoui et al. (2024), and Elliott (2026). While this method is appropriate in some settings, it cannot fit sectors where firms frequently change capacity by substantial and variable amounts.

4 Algorithms and Results

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

uniform investment cost shock on each unit of capital.

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 $k = 1, \ldots, K$, 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.

4.1 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.

To understand the lack of existence, recall that existence of equilibrium in our model relies on the continuity of the mapping $\psi(.)$. 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. 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.[14] For simplicity, we consider uniformly distributed shocks with $\varepsilon_i \sim U(-0.5, 0.5)$.[15] We give the payoff matrix for the modified game below:

		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

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

[15] Unlike our model in Section 2.1, which assumes full support of ε_i over $\Re$, 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.

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.[16]

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 $\pi(2)$.[17] 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

[16] Although one might expect that introducing more draws would reduce the severity of the nonexistence problem, we offer a simple counterexample to this proposition. Consider again the two-draw approximation $\{-0.3, 0.3\}$, for which we showed that the game has a pure-strategy equilibrium. Now add two symmetric interior draws so that the possible draws are$\{-0.3, -0.2, 0.2, 0.3\}$ each with probability (1/4). These additional draws improve the approximation of the underlying $U(-0.5, 0.5)$ distribution: the maximum absolute distance between the empirical and true CDFs decreases from 0.30 to 0.20. Nevertheless, the four-draw game has no pure-strategy equilibrium. Since $p_i \in \left\{0, \frac{1}{4}, \frac{1}{2}, \frac{3}{4}, 1\right\}$, all possibilities can be considered directly. If $p_2 \in \left\{0, \frac{1}{4}\right\}$, Player 1 chooses action (1) for every draw, implying $p_1 = 1$; Player 2 then also chooses action (1) for every draw, implying $p_2 = 1$. If $p_2 = 1/2$, Player 1 chooses action (1) for three of the four draws, so $p_1 = 3/4$; Player 2 then chooses action 1 for every draw, implying $p_2 = 1$. Finally, if $p_2 \in \left\{\frac{3}{4}, 1\right\}$, Player 1 never chooses action 1, so $p_1 = 0$; Player 2 then also never chooses action 1, implying $p_2 = 0$. Thus, each case yields a contradiction.

[17] For ease of notation, we suppress the dependence of π on a_i or i.

remaining active, defined by the solution to the equation:

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

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 the same two values ε_1 and ε_2 for both firms, 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 ($g = 1$).

A symmetric equilibrium does not exist for every value of ε_1 and ε_2. For instance no symmetric equilibrium exists for each of the following four sets of values:

$$\{\varepsilon_1 < \pi(2) < \varepsilon_2 < (1/2)(\pi(1) + \pi(2))\},$$
$$\{\pi(2) < \varepsilon_1, \varepsilon_2 < (1/2)(\pi(1) + \pi(2))\},$$
$$\{(1/2)(\pi(1) + \pi(2)) < \varepsilon_1, \varepsilon_2 < \pi(1)\},$$
$$\{(1/2)(\pi(1) + \pi(2)) < \varepsilon_1 < \pi(1) < \varepsilon_2\}$$

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

4.2 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.[18] 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:

Definition For $1 \leqslant j < k \leqslant K$, let the "$\varepsilon$ *cutoff*" be

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

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

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

[18] Bajari et al. (2007) also use decreasing differences assumption to show monotonicity, referring to it as "Monotone Choice."

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 $\overline{\varepsilon}(1, 2) = 1.5$ and $\overline{\varepsilon}(2, 3) = 0.5$, 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 $0.5 < \varepsilon \leqslant 1.5$, and $a(\varepsilon) = 3$ when $\varepsilon \leqslant 0.5$. Consequently, all three actions are chosen with positive probability.[19] 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: $\overline{\varepsilon}(1, 2) = 1.5 > \overline{\varepsilon}(2, 3) = 0.5$, 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 $\overline{\varepsilon}(1, 2) = 0.5$ and $\overline{\varepsilon}(2, 3) = 1.5$. Also 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.

[19] Note 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

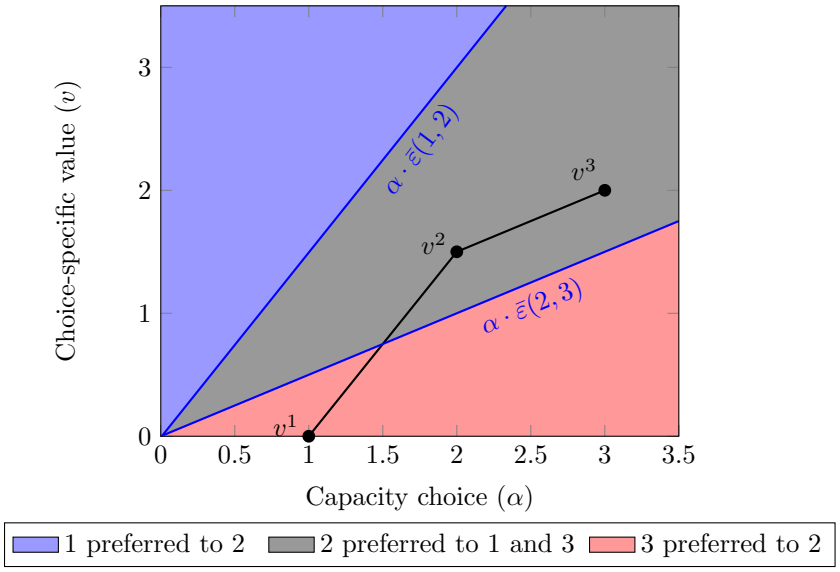
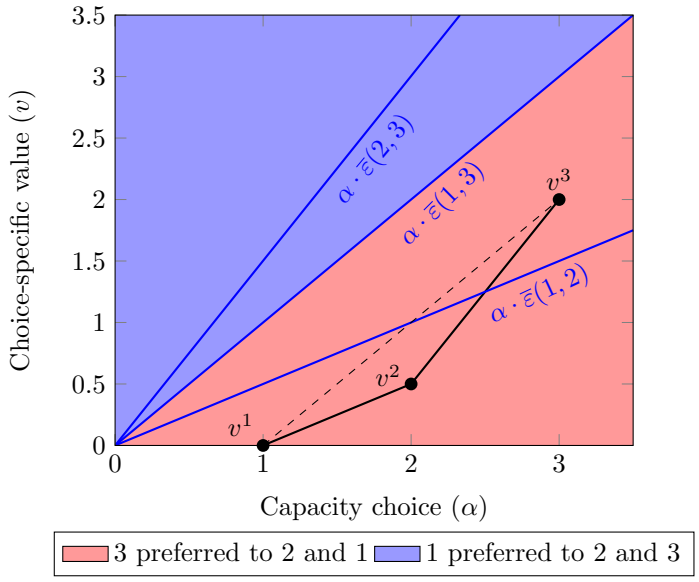


Figure 2: Only two actions chosen with positive probability



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

We now offer our result that characterizes when actions are chosen with positive probability in terms of the ε cutoffs and the discrete convex 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 $\overline{\varepsilon}(k,l) \geqslant \overline{\varepsilon}(j,k)$

(iii) $\forall j, l \in A$ such that $j < k < l$, it holds that $v^k > \lambda v^j + (1-\lambda)v^l$, where $\lambda = \frac{\tilde{c}(\alpha^l) - \tilde{c}(\alpha^k)}{\tilde{c}(\alpha^l) - \tilde{c}(\alpha^j)}$.

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 $\overline{\varepsilon}(k,l) \geqslant \overline{\varepsilon}(j,k)$ then $k \notin C(A)$. If such j, l exist, then for any $\varepsilon > \overline{\varepsilon}(j,k)$, the firm will not pick action k since it prefers action j to action k by Lemma 2. Any $\varepsilon < \overline{\varepsilon}(j,k)$ also satisfies $\varepsilon \leqslant \overline{\varepsilon}(k,l)$, 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 $\overline{\varepsilon}(j,k)$ which has probability 0.

Now, we show that the second statement implies the first. Suppose $\forall j, l \in A$ with $j < k < l$, $\overline{\varepsilon}(k,l) < \overline{\varepsilon}(j,k)$. Now consider $\varepsilon \in [\overline{\varepsilon}^L(k,A), \overline{\varepsilon}^U(k,A)]$, 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.

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

Proposition 1 shows that each action is chosen with positive probability if and only if it is in a discrete analog of the convex 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 $\overline{\varepsilon}^L(j,A) < \overline{\varepsilon}^U(j,A)$. In this case, it is the weakly preferred option when $\varepsilon \in [\overline{\varepsilon}^L(j,A), \overline{\varepsilon}^U(j,A)]$. Finally, $Pr(a = \alpha^j) = Pr(\overline{\varepsilon}^L(j,A) \leqslant \varepsilon < \overline{\varepsilon}^U(j,A))$.*

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

Overall, Proposition 1 and Corollary 1 serve two purposes. First, they characterize when some options are not chosen—and identifies those options that are chosen and those that are not—in terms of the fundamentals of the 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 $\overline{\varepsilon}(j,k)$ for each $1 \leqslant j < k \leqslant K$, 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 growly linearly with K.

4.3 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.[20] 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 convex hull.*
2. *If* $\overline{\varepsilon}(\hat{C}_{k-1}, \hat{C}_k) > \overline{\varepsilon}(\hat{C}_k, \hat{C}_{k+1})$*, then:*[21]
    - *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* $\overline{\varepsilon}(\hat{C}_{k-1}, \hat{C}_k) \leqslant \overline{\varepsilon}(\hat{C}_k, \hat{C}_{k+1})$*, then:*
    - *Drop* $\hat{C}_k$ *so that* $\hat{C} = (\hat{C}_1, \ldots, \hat{C}_{k-1}, \hat{C}_{k+1}, \ldots, \hat{C}_{o(\hat{C})})$.
    - *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:*

[20] Note 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 ordered choices.

[21] We use $\hat{C}_k$ to denote the kth element of $\hat{C}$.

– *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, $j < k < l$, $\overline{\varepsilon}(j,k) \leqslant \overline{\varepsilon}(k,l)$. From Corollary 1, $\overline{\varepsilon}^U(k,A) > \overline{\varepsilon}^L(k,A)$, since $k \in C(A)$. But, by the definitions of the best ε cutoffs, $\overline{\varepsilon}(j,k) \geqslant \overline{\varepsilon}^U(k,A) > \overline{\varepsilon}^L(k,A) \geqslant \overline{\varepsilon}(k,l)$. 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: $\overline{\varepsilon}(\hat{C}_{k-1}, \hat{C}_k) > \overline{\varepsilon}(\hat{C}_k, \hat{C}_{k+1})$ for $k = 2, \ldots, o(\hat{C}) - 1$. Suppose, by contradiction, that $\exists k$ such that $\overline{\varepsilon}(\hat{C}_{k-1}, \hat{C}_k) \leqslant \overline{\varepsilon}(\hat{C}_k, \hat{C}_{k+1})$. 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 $\overline{\varepsilon}(j,k) \leqslant \overline{\varepsilon}(k,l)$. Suppose, again by contradiction, that $\exists k \notin C(A)$ such that $\overline{\varepsilon}(j,k) > \overline{\varepsilon}(k,l), \forall j, l \in C(A)$. Then, consider $\varepsilon \in \left(\overline{\varepsilon}^L(k, C(A)), \overline{\varepsilon}^U(k, C(A))\right)$, 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 $\overline{\varepsilon}(j', \hat{C}_k) \leqslant \overline{\varepsilon}(\hat{C}_k, l')$. 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,

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

We now show a third claim: $\overline{\varepsilon}(\hat{C}_{k-1}, \hat{C}_k) \leqslant \overline{\varepsilon}(\hat{C}_j, \hat{C}_k) \leqslant \overline{\varepsilon}(\hat{C}_j, \hat{C}_{j+1})$. We prove this claim by contradiction. Suppose $\overline{\varepsilon}(\hat{C}_j, \hat{C}_k) < \overline{\varepsilon}(\hat{C}_{k-1}, \hat{C}_k)$. Then, for ε in a neighborhood immediately to the right of $\overline{\varepsilon}(\hat{C}_j, \hat{C}_k)$, $\hat{C}_j$ 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 $\overline{\varepsilon}(\hat{C}_j, \hat{C}_k) > \overline{\varepsilon}(\hat{C}_j, \hat{C}_{j+1})$ there would be an equivalent contradiction.

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

Thus, $\overline{\varepsilon}(\hat{C}_k, l') = \overline{\varepsilon}(\hat{C}_k, \hat{C}_l) < \overline{\varepsilon}(\hat{C}_j, \hat{C}_k) = \overline{\varepsilon}(j', \hat{C}_k)$ which yields a contradiction to our assumption that $\overline{\varepsilon}(j', \hat{C}_k) \leqslant \overline{\varepsilon}(\hat{C}_k, l')$, 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 $1 < k < o(\hat{C})$ for $k \in \hat{C}$ is weakly preferred over all other actions $\iff$ $\varepsilon \in [\overline{\varepsilon}(\hat{C}_k, \hat{C}_{k+1}), \overline{\varepsilon}(\hat{C}_{k-1}, \hat{C}_k)]$, with action 1 preferred for $\varepsilon \in [\overline{\varepsilon}(\hat{C}_1, \hat{C}_2), \infty)$ and action $o(\hat{C})$ preferred for $\varepsilon \in (-\infty, \overline{\varepsilon}(\hat{C}_{o(\hat{C})-1}, \hat{C}_{o(\hat{C})})]$.*

2. *The probabilities of actions are:*

   (a) *Action 1: $Pr\left(\varepsilon > \overline{\varepsilon}(\hat{C}_1, \hat{C}_2)\right)$.*

   (b) *Action $k \in \hat{C}$, $1 < k < o(\hat{C})$: $Pr\left(\overline{\varepsilon}(\hat{C}_k, \hat{C}_{k+1}) \leqslant \varepsilon < \overline{\varepsilon}(\hat{C}_{k-1}, \hat{C}_k)\right)$.*

   (c) *Action $o(\hat{C})$: $Pr\left(\varepsilon \leqslant \overline{\varepsilon}(\hat{C}_{o(\hat{C})-1}, \hat{C}_{o(\hat{C})})\right)$.*

Proof To see that each interior element $\hat{C}_k$ is preferred exactly when $\varepsilon \in [\overline{\varepsilon}(\hat{C}_k, \hat{C}_{k+1}), \overline{\varepsilon}(\hat{C}_{k-1}, \hat{C}_k)]$, 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 $\hat{C}_{o(\hat{C})} = a_{o(\hat{C})}$. 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.[22]

[22]Our algorithm would also likely be faster than a simulation approach. With a simulation approach, one would need to figure out the investment level for each draw. One way of doing this would be to evaluate the value of each investment level for each draw and choose the maximizing level for each draw. With S draws, the computation would be of order $S \times K$. We would expect that a credible application would have $S > K$. With $S = K$, the computation time would be of quadratic order. (One can possibly speed up computation by exploiting Lemma 1, but it is unclear if this would reduce the *order* of computation time.)

5 Simulation Evidence

This section provides simulation evidence on how empirical conclusions depend on the number of investment levels and choice of shock structure, using the context of Gowrisankaran et al. (2026), or GLR for short. GLR model a regulated monopoly utility faced with an energy transition. Each period, the firm decides how much coal capacity to retire and how much combined-cycle natural gas (CCNG) capacity to add. GLR use the model and algorithm of the current paper, or GSD for short. Specifically, the GLR investment and retirement costs both include quadratic mean adjustment costs and an unobservable shock to marginal costs and GLR uses the computation method that we develop here.

This section asks two related questions. First, does it matter whether one uses our proposed modeling approach with a shock to the marginal cost of investment or an alternate specification? Given the preponderance of papers that use *i.i.d.* multinomial logit (MNL) shocks (as noted in Section 1), we focus on this model as the main alternative approach that we investigate. Second, is it appropriate to think of our model as approximating a specification with a large number of closely-spaced investment levels, K?

We provide evidence on these questions from two exercises. First, we estimate different specifications based on GLR that vary in the number of investment levels K and in whether the specification follows GSD or an MNL model with scaled *i.i.d.* shocks on each of the K options. This allows us to understand the impact of model specification on the estimated model parameters. Second, we hold the GLR parameters fixed and examine the impact of progressively finer grids for investment levels, using both GSD and MNL. These simulations isolate the effect of adding investment levels and, in the case of the MNL model, of using independent shocks.

5.1 Parameter estimates across shock specifications

We first investigate using the GLR data but with variation in the number of levels for investment, K, and in whether the costs shocks follow the GSD or MNL model. The GLR model, as described in their equation (4), has one normally distributed unobservable to the

Table 1: Parameter estimates

	GSD				MNL			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Number of grid points:	5	10	20	40	5	10	20	40
Fixed cost of coal	-12.16	0.03	-1.18	-0.16	-52.74	4.92	55.59	147.90
retirement (1e8$)	(11.5)	(8.98)	(3.01)	(5.13)	(549)	(21.7)	(619)	(2541)
Linear coal retirement	1.35	2.41	2.02	2.43	-4.20	2.38	-4.58	-6.52
cost (1e6$/MW)	(1.12)	(1.23)	(0.92)	(1.18)	(59.1)	(0.16)	(104)	(177)
Coal shock std. dev.	0.69	0.08	0.42	0.11	4.94	0.29	6.22	8.74
(GSD:1e6$, MNL:1e9$)	(0.64)	(2.86)	(0.44)	(2.92)	(36.6)	(0.69)	(84.9)	(162)
Fixed cost of CCNG	0.08	-0.12	-0.05	-0.02	3.66	3.07	6.67	7.73
investment (1e8$)	(5.76)	(4.19)	(7.56)	(6.16)	(16.4)	(16.1)	(20.7)	(23.0)
Linear CCNG investment	2.83	2.78	2.82	2.82	2.77	2.73	2.77	2.78
cost (1e6$/MW)	(0.92)	(0.91)	(1.83)	(2.37)	(0.23)	(0.13)	(0.17)	(0.18)
CCNG shock std. dev.	0.00	0.03	0.02	0.01	0.32	0.26	0.29	0.25
(GSD:1e6$, MNL:1e9$)	(0.43)	(0.44)	(0.72)	(0.32)	(0.52)	(0.47)	(0.49)	(0.45)

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. Authors' estimates from data and adjusted specifications of Gowrisankaran et al. (2026). Each column provides results from one specification. Standard errors in parentheses.

marginal cost of adjustment for each technology each period. For ease of exposition, this subsection focuses on the linear model from Table 4 in GLR, rather than their primary specification, which is the quadratic model. In all cases, we use the same nested fixed-point GMM estimation procedure as in GLR. We also use the same moments, which include the difference between the data and the model in the investment and retirement amounts, functions of these amounts, and interactions of these amounts with the existing state.

There are, however, three differences between our specifications and GLR. First, while GLR estimated a model with $K = 20$ capacity levels, we let $K \in \{5, 10, 20, 40\}$ across specifications. Second, while GLR allow for levels to be closer when investment/retirement is less than 1000 MW, we choose equally-spaced levels over the same ranges as GLR to isolate the impact of different values of K. Finally, our MNL specifications assign a scaled *i.i.d.* type 1 extreme value shock to every investment choice.

Table 1 presents estimates from these eight specifications. Starting with column (3), the GSD 20 investment-level model reports parameter estimates that are similar, but not identical, to GLR, due to the fact that we use equally-spaced intervals. Across the three GSD specifications with $K \geqslant 10$ in columns (2)–(4), the parameter estimates are relatively stable. For instance, the linear marginal costs of coal retirement range from 2 to 2.5 and the linear marginal costs of CCNG investment range from 2.75 to 3, both measured in millions of dollars per MW. The sign of the coefficients is the same for all parameters except for the fixed costs of coal retirement (which is never significant, mirroring GLR). The specification with 5 levels is more different, with a very large but insignificant fixed cost of coal retirement. This suggests that using 10 or more levels is important.

In contrast, the MNL model has parameter estimates that are less stable, particularly for coal retirement. For instance, the coal linear marginal cost coefficient flips signs twice in the MNL model, when going from 5 to 10 to 20 investment levels, in columns (5) to (7). Thus, the economic implications of the MNL model for this context would be very different than from the GSD model. Interestingly, the estimated MNL shock scales do not decline systematically as K increases. The CCNG shock standard deviation remains between 0.25 and 0.32, while the coal shock standard deviation is erratic and very imprecisely estimated.

Combined with the fact that there are K *i.i.d.* draws, this implies that this model will place more weight on unobservables as the number of investment levels increases.

Overall, for this empirical context, once a minimum number of investment levels has been reached, estimates from a model with one shock to marginal cost are more robust to the number of investment levels than a model with an *i.i.d.* shock to every available choice. This likely reflects a basic distinction between the specifications. Adding investment levels in the GSD model gives the firm a finer approximation to the same ordered adjustment problem. In the MNL model, every additional level changes both the action set and the stochastic environment, by giving the firm an additional independent shock. The parameter estimates will then adjust to account for both these sources of variation.

5.2 Investment distributions across shock specifications

Given our finding that estimated parameters vary substantially across specifications and across the number of investment levels for the MNL model, we next hold the structural parameters fixed to isolate the effect of the shock specification and investment level refinement on model predictions. For this exercise, we use the GLR sample and principal specification, which has quadratic terms on the investment/retirement costs, allowing for $K = 20$ investment levels as in GLR. However, as in Section 5.1, we let those investment levels be equally spaced and hence we reestimate the parameters, obtaining values similar to Table 4 in GLR. An important question is how to scale the MNL residuals. In GLR, the residual is normal and is scaled by estimated parameters σ^f, for fuel/technology type $f \in \{CCNG, coal\}$. While there is no way to exactly replicate the variation in the marginal cost shock with *i.i.d.* shocks, we scale each *i.i.d.* shock by the standard deviation of the difference between investment levels,

$$\sigma^{f,MNL} = 157.9 \times \sqrt{3}/\pi \times \sigma^f, \tag{10}$$

using 263.1 for coal. In (10), the third term is the same as in GLR, the second term scales divides by the standard deviation of the type 1 extreme value distribution, and the first term is the difference between subsequent investment levels.

Using these parameters, we first compute the Bellman equations across values of K and with the GSD and MNL unobservable specifications. For each utility, we then simulate optimal dynamic decisions for 10 periods, using as its starting state its coal and NGCC capacity at the first period it is observed in the GLR data the and the 2007-09 natural gas fuel price. We repeat this simulation exercise 10,000 times and report statistics on initial expected discounted value and investment across the 10 periods. Although the model includes both coal retirement and CCNG investment, this subsection reports capacity investment only.

Table 2 presents base simulation results for both models. Considering first the mean and standard deviation of investment in columns (1) and (3), we find that our model converges to very similar values for 20 or more investment choices, with a mean investment level of about 282 MWs per period and a standard deviation of about 626 MWs. Recall that GLR chooses discrete investment levels as an approximation to a continuous choice of investment. The fact that the discrete approximation appears to converge provides reassurance that the discrete choice model with shocks to the marginal cost of investment and discrete investment levels approximates a model with many ordered choices.

Turning to the model with multinomial logit shocks in columns (2) and (4), mean investment continues to increase in K and the standard deviation of investment decreases. 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.

Columns (5) and (6) show the mean initial expected discounted values for the two models. The model with shocks to the marginal cost of investment finds mean values that converge to about \$2.8 billion as the number of choices increases. As with investment decisions, we find little variation with more than 20 investment levels. 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 confirm that the values continue to grow as the number of choices increases.

One might also view the multinomial logit model as an approximation to a continuous

Table 2: Base simulations

	Mean investment (MW)		Std. dev. investment (MW)		Mean value (millions of 2006 $)			Mean # zero prob. choices
	GSD	MNL	GSD	MNL	GSD	MNL	AR	GSD
# inv. choices	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
5	257.1	91.4	630.3	322.7	2,643	2,408	2,184	0.0
10	274.3	146.8	628.3	325.8	2,694	2,460	2,120	0.0
20	281.1	185.0	627.0	322.9	2,723	2,585	2,114	0.0
30	282.6	194.1	627.9	319.2	2,738	2,674	2,118	0.1
40	283.3	197.4	627.6	317.6	2,749	2,741	2,123	0.1
50	282.9	199.4	626.4	315.4	2,759	2,794	2,125	0.1
60	281.9	200.3	625.5	315.1	2,769	2,837	2,128	0.1
70	282.9	200.8	626.3	314.0	2,777	2,874	2,129	0.2
80	282.4	201.1	626.1	313.6	2,785	2,905	2,130	0.2
90	281.6	201.4	624.8	313.2	2,792	2,934	2,131	0.2
100	282.0	202.0	625.9	312.4	2,798	2,959	2,132	0.2

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 Ackerberg and Rysman (2005) correction. Computations from GLR sample and parameters with 10,000 simulation draws.

investment model, where the standard deviation of the multinomial shock is a tuning parameter that would decline as K increases. Indeed, Ackerberg and Rysman (2005) suggest a correction to the static multinomial model with large numbers of options that, applied to our context, implies adjusting the value by subtracting $\sigma^{f,MNL} \times \log(K)$ from the Bellman equation. Although the structural interpretation of their correction is not applicable to settings with capacity choices, column (7) nonetheless computes a variant of our model with the correction. It shows that, in this case, values appear to converge as K increases.

Finally, we investigate whether all investment levels are chosen with positive probability in our model. Column (8) shows the mean number of levels that are never chosen for sample utility/period observations in the GLR sample.[23] In the base model, on average, 0.1 levels are never chosen with 40 investment levels. This statistic increases to 0.2 when allowing for 100 investment levels. Thus, this issue appears to be present but very limited for this specification.

Some of the investment results in Table 2 may be specific to the fact that 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. Counterfactual policies such as license fees could also increase fixed costs of investment.

To understand the implications of fixed costs on these results, Table 3 provides simulations from the same model, but setting the fixed cost parameter of coal retirement or CCNG investment to $500 million, which is roughly the estimated linear mean marginal costs of a small 200 MW generator.

As with the base specification, we find that our model is stable with 20 or more choices. Given the higher fixed costs of CCNG investment, it is not surprising that our model predicts a lower level of investment than with the Table 2, with mean investment per period dropping from roughly 282 MWs to 207 MWs. The standard deviation of investment rises slightly.

The MNL model also has lower mean investment. As with the base simulations, mean investment does not converge as K increases. The Ackerberg and Rysman (2005) model

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

Table 3: Simulations with significant fixed costs

	Mean investment (MW)		Std. dev. investment (MW)		Mean value (millions of 2006 $)			Mean # zero prob. choices
	GSD	MNL	GSD	MNL	GSD	MNL	AR	GSD
# inv. choices	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
5	187.2	35.0	650.7	284.2	2,559	2,395	2,168	0.9
10	199.4	36.0	664.6	281.9	2,577	2,403	2,030	3.3
20	205.0	38.3	670.0	284.2	2,586	2,427	1,909	8.2
30	206.2	40.6	672.1	287.1	2,587	2,451	1,852	13.4
40	206.7	43.5	671.4	292.5	2,588	2,472	1,817	18.3
50	208.1	46.1	674.1	297.4	2,589	2,491	1,790	23.5
60	207.0	48.7	672.4	301.4	2,589	2,508	1,773	28.5
70	206.8	51.3	672.6	305.3	2,589	2,524	1,759	33.6
80	207.9	53.4	674.1	308.3	2,590	2,539	1,748	38.5
90	206.6	55.5	671.8	313.1	2,590	2,553	1,739	43.6
100	206.5	57.6	671.6	315.8	2,590	2,565	1,730	48.7

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 Ackerberg and Rysman (2005) 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.

now overcorrects the Bellman equation substantially, with the *ex ante* value dropping by 9% when going from 20 to 100 investment levels. The difference between these results and those in Table 2 shows that it may be difficult to find a tuning parameter that would consistently yield stable values across contexts.

The differences in the relative implications of the MNL model across Tables 2 and 3 also suggests that the MNL model may provide misleading counterfactual policy implications in cases where the data reflect shocks to firms' marginal cost of investment.

Importantly, the number of investment levels that are never chosen is dramatically higher than in Table 2 and increases more than proportionally in K. With $K \geqslant 20$, it now includes a third or more of investment levels. This occurs because, in many cases, 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.

The substantial share of options never chosen in this case illustrates that fixed costs are an economic force that creates non-concave choice-specific value functions. This non-concavity then results in many choices not being on the discrete convex hull. As noted in Section 1, this non-concavity also suggests that one might use a nested fixed-point GMM estimation algorithm in these settings, because it accommodates the possibility of zero probability choices.

Taken together, the two exercises provide complementary evidence. The estimation results show that changing the grid can alter economically important cost estimates under independent choice-specific shocks, especially for coal retirement. The fixed-parameter simulations show the mechanism: adding choices under MNL also adds shocks and changes predicted investment and *ex ante* values. In contrast, GSD yields stable predictions as the grid is refined while allowing feasible choices to have zero probability.

6 Conclusion

This paper develops a new method that is useful in computing and estimating dynamic models where firms invest to build or retire capacity, and where the potential investment

levels have many ordered values. This class of models has been used repeatedly in the recent literature, in settings ranging from electricity utilities building new generators to shipbuilders adding to their order queue. Our method allows for a shock structure that is useful in many settings and can be used even in the presence of non-concave investment responses, which can occur with investment fixed costs or oligopoly interactions.

In general, having unobservable shocks to the cost of investment is useful in this setting for ensuring existence of pure strategy equilibrium, generating variation in predicted choices, and ensuring a smooth GMM or likelihood criterion function in estimation. We argue that shocks to the marginal cost of investment are more economically intuitive than other types of shocks, resulting in them being used by an important subset of papers within this class of models. Section 5 also shows that this type of shock is more robust than other specifications as we vary K, the number of potential investment levels.

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 convex hull of the choice-specific value functions relative to other actions. In many important economic cases, models with *i.i.d.* shocks to the marginal costs of investment will have some investment levels never chosen due to non-concavity of the choice-specific value functions. For instance, Section 5.2 provides simulation evidence that shows that this will occur frequently in the presence of substantial fixed costs of investment.

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, while a naïve algorithm would require $K(K-1)/2$ comparisons, and hence is slower for $K > 3$.

We offer the following advice to practitioners. First, we recommend that researchers use private information *i.i.d.* shocks to the marginal cost of investment in settings that require unobservables, which includes most dynamic estimation and dynamic oligopoly computation exercises. Second, we recommend that researchers use our method in cases where there the choice-specific value functions may not be concave. This includes models with fixed costs of

investment and increasing returns to investment, and dynamic oligopoly models. Our method can be used in many types of conditional choice probability estimation and in equilibrium computation. Equilibrium computation is useful for nested fixed-point estimation—which we recommend implementing with a GMM estimator—and in counterfactual computations.

Nonetheless, we note two important limitations of our algorithm. First, while our method improves the speed of each step of the successive approximations algorithm and may increase the speed of forward simulation since not every state is reached in equilibrium, it does not address the computational problem of a curse of dimensionality arising from large state spaces (Rust, 1997). Second, one-step ahead estimators (Arcidiacono and Miller, 2011) cannot directly be used with our model, because the presence of zero probability choices implies that they do not generate an *invertible* mapping from values to probabilities. It may be possible for future research to develop inequality-based estimators that adapt our approach to this type of estimator.

References

Ackerberg, 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.

Arcidiacono, P. and Miller, R. A. (2011). Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica*, 79(6):1823–1867.

Bajari, P., Benkard, C. L., and Levin, J. (2007). Estimating dynamic models of imperfect competition. *Econometrica*, 75(5):1131–71.

Barwick, P. J., Kalouptsidi, M., and Zahur, N. B. (2025). Industrial policy implementation: Empirical evidence from China's shipbuilding industry. *Review of Economic Studies*, 92(6):3611–3648.

Benkard, C. L. (2004). A dynamic analysis of the market for wide-bodied commercial aircraft. *The Review of Economic Studies*, 71(3):581–611.

Bresnahan, T. F. (1987). Competition and collusion in the american automobile industry: The 1955 price war. *Journal of Industrial Economics*, 35(4):457–482.

Caoui, E. H. (2023). Estimating the costs of standardization: Evidence from the movie industry. *The Review of Economic Studies*, 90(2):597–633.

Caoui, E. H., Hollenbeck, B., and Osborne, M. (2024). Dynamic entry & spatial competition: An application to dollar store expansion. *Available at SSRN 4872647*.

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 Satterthwaite, M. (2010). Computable Markov-perfect industry dynamics. *The RAND Journal of Economics*, 41(2):215–243.

Elliott, J. T. (2026). Investment, emissions, and reliability in electricity markets. *American Economic Review*, forthcoming.

Ericson, R. and Pakes, A. (1995). Markov-perfect industry dynamics: A framework for empirical work. *The Review of Economic Studies*, 62(1):53–82.

Fowlie, M., Reguant, M., and Ryan, S. P. (2016). Market-based emissions regulation and industry dynamics. *Journal of Political Economy*, 124(1):249–302.

Gowrisankaran, G. (1995). A dynamic analysis of mergers. Doctoral Dissertation, Yale University.

Gowrisankaran, G., Langer, A., and Reguant, M. (2026). Energy transitions in regulated markets. *American Economic Review*, 116(8):2928–2961.

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.

Hotz, V. J. and Miller, R. (1993). Conditional choice probabilities and the estimation of dynamic models. *Review of Economic Studies*, 60:497–529.

Igami, M. and Uetake, K. (2020). Mergers, innovation, and entry-exit dynamics: Consolidation of the hard disk drive industry, 1996–2016. *The Review of Economic Studies*, 87(6):2672–2702.

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.

Kalouptsidi, 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. 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.

Pakes, A., Ostrovsky, M., and Berry, S. (2007). Simple estimators for the parameters of discrete dynamic games. *RAND Journal of Economics*, 75(5):1331.

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.

Wang, Q. (2026). Innovation path choices in China's electric vehicle battery industry. *Available at SSRN 6168655*.

Yun, Y. (2025). Navigating the commons.