Bounding Outcomes in Counterfactual Analysis*

Mar Reguant
Northwestern University, BSE, CEPR and NBER

Francisco Pareschi
Northwestern University

July 2023

Abstract

In many economic settings, counterfactual analysis can be difficult for two reasons: (i) we do not know how to compute the equilibrium of the game, or (ii) even if we know how to compute one equilibrium, the game might feature multiple equilibria, which are challenging to characterize exhaustively. We propose a bounding framework to allow for counterfactual analysis even when these problems might arise. The method relies on determining valid (conservative) bounds to counterfactual outcomes that contain any outcome that could be sustained in equilibrium, i.e., any outcome that can be supported by a set of equilibrium constraints. To ensure that all potential solutions are considered, We propose to reframe equilibrium constraints as a relaxed mixed-integer linear program. We show that the framework can also be used to narrow down equilibria by imposing additional equilibrium constraints. We provide examples of static price competition with differentiated products, dynamic games, and multi-unit auctions, three areas where counterfactual analysis faces these challenges.

*Email: mar.reguant@northwestern.edu. We thank Pierre Dubois, Gautam Gowrisankaran, Liran Einav, Sébastien Houde, Marti Mestieri, Rob Porter, and Peter Reiss for useful discussions.

1 Introduction

In many economic settings, counterfactual analysis can be difficult for two reasons: (i) we do not know how to compute the equilibrium of the game, or (ii) even if we know how to compute one equilibrium, the game might feature multiple equilibria, which are challenging to characterize exhaustively. We propose a new methodology to do counterfactual analysis even when these problems might arise. The method relies on determining valid (conservative) bounds to counterfactual outcomes that contain any outcome that could be sustained in equilibrium (and possibly more). Because the bounds are conservative, they are easier to compute than tight bounds, albeit less informative.

The motivation for the method is based on common challenges faced by researchers in many branches of applied economics. One of the main goals of empirical work, and especially structural modeling, is to perform counterfactual analyses to predict the impact of policies out of sample. To do so, the researcher first poses a model, which is estimated and then used to perform counterfactual experiments. Whereas these steps are conceptually straightforward, the researcher faces several challenges. First, the empirical model needs to be sufficiently rich to capture the observed patterns in the data meaningfully. At the same time, the solution to the posed empirical model needs to be tractable enough. In many situations, it is hard to characterize the solution to the structural model, making estimation and/or computation of counterfactuals difficult. Even if one can characterize one equilibrium, several equilibria might exist.

Estimation methods can often circumvent some of the issues of equilibrium calculation by applying indirect inference arguments. For example, in the context of dynamic games, one can estimate transition probabilities, or expected continuation values, from the policies observed in the data (Bajari et al., 2007; Pakes et al., 2007).[1] However, such approaches do not work when computing counterfactuals, as those are typically about experiments out of sample. Whereas one can often find one equilibrium in the game, exhaustively characterizing multiple equilibria can be difficult (Besanko et al., 2010). Similar issues arise in the literature on empirical auctions, e.g., in multi-unit auctions or auctions with minimal behavioral assumptions. Necessary and/or sufficient conditions are commonly used to estimate or bound the underlying valuations of the bidders in a given auction or set of auctions (Guerre et al., 2000; Haile and Tamer, 2003; Jofre-Bonet and Pesendorfer, 2003; Kastl, 2011; Hortaçsu and McAdams, 2010). However, it is often difficult to compute the equilibrium of such auctions. Indeed, we are often limited in the type of counterfactuals we can consider, e.g., to mechanisms that are truthful (Kastl, 2011; Hortaçsu and McAdams, 2010; Kang and Puller, 2008). Even in relatively straightforward games of price competition with differentiated products, there is no guarantee that a unique equilibrium exists.

We present a new method to bound counterfactual outcomes under relatively general conditions. The basic idea of the method is intuitive. Imagine we are interested in finding bounds to the welfare that would arise after a policy change. For example, in the context of merger simulations, this could be the expected welfare changes that could occur after a merger. In the context of environmental policy, we could be interested in quantifying welfare bounds after implementing a carbon tax. The proposed method searches for an upper and lower bound to a particular counterfactual outcome (e.g., welfare), subject to the outcome

[1] The issue of multiple equilibria in estimation is still an active area of research.

being implementable in equilibrium. Therefore, it involves minimizing and maximizing a given equilibrium outcome of interest (welfare) subject to equilibrium constraints.[2]

Define an equilibrium outcome of interest as $W(\mathbf{x};\theta)$, where $\mathbf{x}$ represents the strategy space of the players and θ are the fundamental parameters of the model, which are known at this stage. We propose to search for the minimum and maximum of a counterfactual outcome W subject to the equilibrium constraints of the game, i.e.,

$$\begin{aligned} \underline{W}(\theta) \equiv \min_{\mathbf{x}} \quad & W(\mathbf{x};\theta) \\ \text{s.t.} \quad & \mathbf{G}(\mathbf{x};\theta) = 0, \end{aligned} \tag{1}$$

and,

$$\begin{aligned} \overline{W}(\theta) \equiv \max_{\mathbf{x}} \quad & W(\mathbf{x};\theta) \\ \text{s.t.} \quad & \mathbf{G}(\mathbf{x};\theta) = 0, \end{aligned} \tag{2}$$

where $\mathbf{G}(\mathbf{x};\theta)$ is a system of equations (it could also be inequalities) that describes the equilibrium conditions of the game. Note that the method only searches for the minimum and maximum outcomes that can be achieved when equilibrium conditions are satisfied. This contrasts with homotopy methods, which try to characterize as many equilibria as possible, but it has the advantage of being more computationally tractable.

The researcher can compute bounds on a given outcome to perform counterfactual experiments as the fundamental parameters change. Consider a policy change θ'. Instead of examining how point predictions change at a given equilibrium, the proposed method compares outcome bounds under the original policy, i.e., $[\underline{W}(\theta), \overline{W}(\theta)]$, to counterfactual bounds under the new policy, i.e., $[\underline{W}(\theta'), \overline{W}(\theta')]$. If the two sets are non-overlapping, then one can sign the effects of a counterfactual policy change. If not, then the predictions are more nuanced and will depend on which equilibrium is being played.

To fix ideas, consider a simple static entry game with multiple equilibria.[3] The example is inspired by Fowlie et al. (2014), who consider entry and exit decisions by cement plants. Suppose there is a dirty firm with a large capacity and a clean firm with a smaller capacity competing in quantities. They have to pay a fixed operating cost to stay in the market. If a carbon tax is imposed, firms might exit the market. If only one of them exits the market, which of the two will? What happens as we change the carbon tax? Figure 1 shows that multiple equilibria might arise depending on the carbon tax. In particular, for low to medium ranges of carbon taxes (between \$5 and \$22, approximately), either the dirty or the clean plant might exit the market, leading to substantially different emissions reductions.

Figure 1 also shows the equilibrium emissions from a policy that introduces an output subsidy (dashed line). Because the output subsidy is not contingent on emissions levels, it benefits the clean plant relatively more. In such a case, the outcomes under the carbon tax plus a subsidy are more tilted towards the lower emissions equilibria for low carbon prices. For example, a carbon tax of \$20 dollars plus a \$16 output

[2] This is what is often known as a mathematical program subject to equilibrium constraints (MPEC).

[3] Multiple equilibria is common in entry/exit games of complete information (Bresnahan and Reiss, 1990).

Figure 1: A simple example of counterfactual bounds

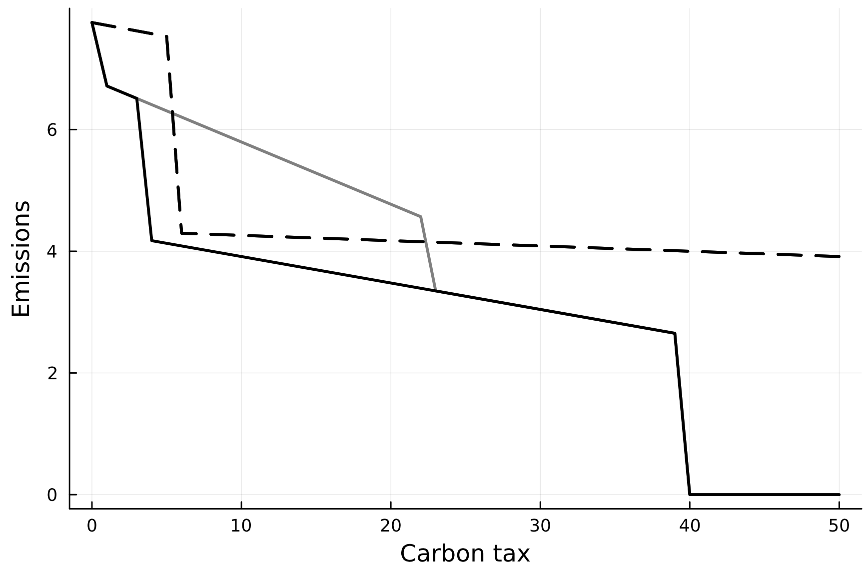

A simple entry game in which there is a dirty firm with quadratic costs $C(Q_1) = 20Q_1 + Q_1^2 + 1.5\theta Q_1$, and a clean firm with quadratic costs $C(Q_2) = 20Q_2 + 1.5Q_2^2 + \theta Q_2$, where θ represents the carbon tax. Firms maximize profits a la Cournot with complete information on the firms that remain in the market. The fixed costs of operating are 100 and 80, respectively. The inverse demand is given by P = 120 - 10 Q. In equilibrium, the profit of firms remaining active needs to at least cover the fixed cost.

subsidy (80% of the tax burden at the efficient emissions rate of 1.0) achieves the same equilibrium emissions as the lower bound when the carbon tax is \$10. Still, it does not support an equilibrium in which the dirty plant stays.

This example illustrates the importance of accounting for the multiplicity of equilibria when modeling counterfactual experiments. In this stylized setting, it is easy to check for all equilibria exhaustively. One can find the lowest and highest levels of counterfactual outcomes more generally by finding their minimum and maximum (in this case, emissions) subject to the participation constraints of the firms and their optimal quantity decisions being satisfied. The advantage of searching for counterfactual bounds directly is that only those equilibria that give the most extreme equilibrium outcomes are fully characterized, which can be particularly helpful in richer environments. Whereas this is often less informative than the full set of equilibria, it is easier to compute. Furthermore, the proposed counterfactual bounds are valid and, thus, enable making decisions that are robust to multiple equilibria.

From a technical point of view, the search for counterfactual bounds established by (1) and (2) can be implemented with a non-linear constrained solver. The solver finds lower and upper bounds to the outcome W under alternative scenarios θ and θ'. If the maximization and minimization problems are well-behaved, one can ensure that valid counterfactual bounds to equilibrium outcomes have been characterized. In general, however, non-linear solvers are not robust. i.e., they might not find the global minimum and maximum. This limitation is often shared with homotopy methods, which can characterize many equilibria but not necessarily all of them.[4]

To circumvent this problem, we convert the above program into a relaxed mixed-integer piece-wise linear program by creating a piece-wise linear envelope around the non-linear constraints.[5] This presents several advantages. First, the linearity and the integer variables ensure the solution is a global optimum. Second, because it is a relaxed version of the original problem, it necessarily includes the true bounds of the objective function. Therefore, even though the method relies on approximation techniques, it does not compromise the validity of the bounds. We show that the particular approximation used to compute the bounds does not affect the validity of the bounds, although the bounds themselves might be conservative. We also propose refinement techniques to increase the sharpness of the bounds in practice.

The method can be applied to a variety of settings. In particular, we show how to use it to compute robust bounds to counterfactual outcomes as long as the strategy space is finite and bounded in equilibrium. It also requires that the fundamentals of the problem (e.g., a known cost function) can be bounded in some meaningful way.[6]

To demonstrate the applicability of the proposed methodology, we present three typical applications in Industrial Organization, in which equilibrium computation has been challenging. First, we illustrate the applicability of our method to pricing games by oligopolistic firms facing discrete choice demand systems. Second, we show how to apply the method to dynamic games, using a model of learning and forgetting that builds on Ericson and Pakes (1995). This is a useful example, as the game was previously identified as susceptible to multiple equilibria using homotopy techniques (Besanko et al., 2010). This example also highlights how the proposed methodology can be used as a powerful equilibrium refinement tool. Finally, we show how the methodology can compute bounds to equilibria in multi-unit auctions, when firms compete in price-quantity schedules.

Related Literature. This paper is related to the literature examining how to compute equilibria in games, especially in the presence of multiple equilibria. Therefore, it is related to homotopy methods (Besanko et al., 2010), which have also been proposed to perform counterfactual analysis (Aguirregabiria, 2012). It is also related to the literature exploiting recent algorithms to compute all equilibria in discrete games of complete information (Bajari et al., 2010). Aguirregabiria and Mira (2012) propose to use a genetic algorithm coupled with a nested fixed point algorithm to explore the possibility of multiple equilibria. Instead of all (or many) equilibria, we propose to bound the counterfactual outcome of interest conservatively. Our methods are based on optimization approaches that rely on mixed-integer linear formulations and are more broadly related to constrained optimization approaches (Su and Judd, 2012; Dubé et al., 2012).

There is a growing literature on how to bound counterfactual outcomes in a robust manner.[7] For example,

[4] This is not always the case. See Judd et al. (2012) for an example in which *all* equilibria can be computed. Bajari et al. (2010) also shows how to compute all equilibria in discrete games of complete information.

[5] This is the essence behind deterministic global optimization methods, which directly used to solve for the bounds. However, there are limits to the class of mathematical problems that some solvers can consider.

[6] This might seem restrictive at first sight. However, even though costs might be infinity for some quantity choices, it is only required that they can be bounded within the "relevant" range. There are natural ways to narrow down the relevant range of approximation. For example, we know it is not optimal for firms to produce in regions with infinity marginal costs. We also show how to narrow down this relevant range using the proposed methodology.

[7] There is an extensive literature on how to bound parameters in the estimation context. We focus here on papers that emphasize

Jia (2008) exploits supermodularity in an entry game to focus on the highest profit equilibria, enabling parameter estimation. Uetake and Watanabe (2014) also studies an entry game with a lattice structure and computes minimum and maximum counterfactual outcomes considering that the parameters of the model are only set identified. Sweeting (2009) considers multiple equilibria when studying commercial timing in radio stations by exploring the best responses starting at extreme points and gradually moving along them. Grennan and Town (2015) examines an innovation game in which bounds to social surplus are computed from a set of necessary equilibrium conditions that are easier to compute than the full equilibrium.

Finally, the examples presented are related to the literature on supply function equilibria and dynamic games. We discuss the relevant literature in more detail in the context of each particular application.

2 Methodology

We propose practical methods to find maximum and minimum bounds to a counterfactual outcome of interest, $W(\mathbf{x};\theta)$, where $\mathbf{x}$ represents the vector of equilibrium strategies, and θ are the fundamental parameters of the model, which are known to (or have been estimated by) the researcher.

The three pillars of the method are *non-iteration*, *piece-wise linear approximation*, and, more broadly, *relaxation*. Before getting into these concepts, it is important to state the following assumptions.

Assumption 1. $\mathbf{x} \in [\underline{\mathbf{x}}, \overline{\mathbf{x}}]$.

Assumption 2. $\dim\{\mathbf{x}\} < \infty$.

Assumption 1 ensures that one can limit the relevant range of values that strategic variables can take. This assumption is less restrictive than it might seem. For example, take a pricing game. Whereas prices could indeed take any real value (e.g., they could potentially be negative or infinity), in practice, such ranges are not relevant for any equilibrium. Therefore, the relevant assumption is that such variables can be bounded in equilibrium.

Assumption 2 limits the size of the strategy space to be finite. This could be relevant in continuous games, e.g., when agents choose contingent plans over continuous states. To palliate this limitation, we can use approximation techniques to transform a game with continuous strategies to one that depends on a limited set of variables, e.g., where firms choose parameters of a flexible function. An alternative would be to discretize the number of states.

2.1 Non-iteration

One building block of the methodology is to express the problem as a one-step optimization problem. A one-step program has the advantage of avoiding iteration methods (e.g., Gauss-Jacobi or Gauss-Seidel). This is crucial to avoid "missing" particular equilibria in a game, as iteration methods might converge to specific regions of the strategy space (e.g., iteration methods might tend to converge to stable equilibria).

equilibrium and counterfactual computation. The current discussion is not intended to be exhaustive.

The constraints of the program are the equilibrium constraints of the game.[8] Importantly, the objective function is not necessary to define the equilibrium constraints of the game. Instead, the objective function is reserved for the counterfactual outcome the researcher is interested in bounding. This is a crucial feature, as one can directly minimize and maximize the objective function to characterize the bounds.[9] Having the counterfactual outcome as the objective function also has the advantage that the algorithm does not try to characterize all the solutions that could satisfy the system of equilibrium constraints. Instead, it searches for the maximum and minimum counterfactual outcomes that can be sustained in equilibrium.

Proposition 1. *Assume that one has necessary and sufficient conditions to a game in the form of piece-wise linear constraints, defined by a vector* $\mathbf{G}(\mathbf{x};\theta) = 0$*. Assume that we can define a counterfactual outcome that is piece-wise linear,* $W(\mathbf{x};\theta)$*.*

Then, the solution to

$$\begin{aligned} \overline{W}(\theta) \equiv \max_{\mathbf{x}} \quad & W(\mathbf{x};\theta) \\ \textit{s.t.} \quad & \mathbf{G}(\mathbf{x};\theta) = 0, \end{aligned}$$

and,

$$\begin{aligned} \underline{W}(\theta) \equiv \min_{\mathbf{x}} \quad & W(\mathbf{x};\theta) \\ \textit{s.t.} \quad & \mathbf{G}(\mathbf{x};\theta) = 0, \end{aligned}$$

provides valid sharp bounds to the counterfactual outcome W*, where the bounds are given by* $[\underline{W}(\theta), \overline{W}(\theta)]$*.*

Proof. The problem is defined as a mixed integer linear program (MILP). MILP can be solved exhaustively in a robust manner (up to computational limitations). The bounds are sharp because $\mathbf{G}(\mathbf{x};\theta)$ represents a set of necessary *and sufficient* conditions to equilibrium. $\square$

This representation is also useful as a device to detect multiple equilibria. If there is a single equilibrium that satisfies $\mathbf{G}(\mathbf{x};\theta) = 0$, then $\underline{W}(\mathbf{x};\theta) = \overline{W}(\mathbf{x};\theta)$.[10] If $\underline{W}(\mathbf{x};\theta) < \overline{W}(\mathbf{x};\theta)$, then there exist multiple equilibria.

One could be worried that, whereas mixed integer linear programs can be solved exhaustively, this is not possible in practice.[11] However, current algorithms have improved substantially, making previously intractable problems within reach. A class of algorithms that has been particularly successful is the so-called "branch-and-bound" algorithms. These algorithms are beneficial for computing conservative bounds that account for the computational burden of the program.

[8] This is often known as a mathematical program subject to equilibrium constraints (MPEC).

[9] Similar constrained approaches have also been used in estimation. In that case, the likelihood function constitutes the objective functions, also subject to equilibrium constraints (Dubé et al., 2012).

[10] Note that this is a necessary condition for the equilibrium to be unique, but not sufficient, as alternative combinations of strategies might lead to the same counterfactual value.

[11] In optimization and computer science, mixed integer linear programs are known to be "NP-hard" (non-deterministic polynomial-time hard), which implies that an algorithm that is polynomial might not exist.

Even when the algorithm fails to explore all potential solutions exhaustively, branch-and-bound algorithms provide an upper and lower bound to the objective function (the "optimality gap"), which can be used to ensure that the counterfactual bounds are valid. Importantly, this requires that at least one solution can be found.

Observation 1. *As long as one feasible solution can be characterized, the optimality gap provided by branch-and-bound algorithms offers upper and lower bounds that account for computational limitations in the solution. One can use the optimality gap to ensure that bounds are conservatively valid.*

2.2 Piece-wise Linear Approximation

The optimization approach outlined above only ensures that valid counterfactual bounds can be found for games where the objective function and constraints are piece-wise linear, which guarantees that a global minimum and maximum can be exhaustively searched for. However, the nature of most economic problems is far from linear. Can one still use the above method?

As a second building block, we propose to approximate the equilibrium conditions of the game with piece-wise linear envelopes.[12] The linearity ensures that all solutions can be considered, even in the presence of nonlinear constraints that would otherwise be difficult to explore. However, it comes at the expense of losing the sharpness of the bounds.

Proposition 2. *Assume that one can define necessary and sufficient conditions to a game in the form of a vector $\mathbf{G}(\mathbf{x};\theta) = 0$, which one can bound over $[\underline{\mathbf{x}}, \overline{\mathbf{x}}]$ with piece-wise linear approximations, $\underline{\mathbf{G}}(\mathbf{x}, \mathbf{u}; \theta) \leq \mathbf{G}(\mathbf{x};\theta) \leq \overline{\mathbf{G}}(\mathbf{x}, \mathbf{u}; \theta)$, where $\mathbf{u}$ is a set of auxiliary integer variables, $\mathbf{u} \in \{0, 1\}$. Assume that we can define a counterfactual outcome that is piece-wise linear in equilibrium strategies, defined by $W(\mathbf{x}, \theta)$.*[13]

Then, the solution to

$$\begin{aligned} \overline{W}_u(\theta) \equiv \max_{\mathbf{x}} \quad & W(\mathbf{x};\theta) \\ \textit{s.t.} \quad & \underline{\mathbf{G}}(\mathbf{x}, \mathbf{u}, \theta) \leq 0, \\ & \overline{\mathbf{G}}(\mathbf{x}, \mathbf{u}, \theta) \geq 0, \end{aligned}$$

and,

$$\begin{aligned} \underline{W}_u(\theta) \equiv \min_{\mathbf{x}} \quad & W(\mathbf{x};\theta) \\ \textit{s.t.} \quad & \underline{\mathbf{G}}(\mathbf{x}, \mathbf{u}, \theta) \leq 0, \\ & \overline{\mathbf{G}}(\mathbf{x}, \mathbf{u}, \theta) \geq 0, \end{aligned}$$

provides valid bounds to the counterfactual outcome W, where the bounds are given by $[\underline{W}_u(\theta), \overline{W}_u(\theta)]$.

[12] Similar arguments can be used to relax the assumption that the objective function is piece-wise linear, as explained below.

[13] Note that this setting can easily accommodate the case in which the objective function is nonlinear. Without loss of generality, one can always define the objective function as an auxiliary variable, making the objective function linear. The definition (and envelopes) to such auxiliary variable become part of the constraints in G.

These bounds are (weakly) conservative with respect to sharp bounds, i.e., $\underline{W}_u(\theta) \leq \underline{W}(\theta)$*, and* $\overline{W}_u(\theta) \geq \overline{W}(\theta)$.

Proof. This is a particular case of a relaxed optimization problem. The optimum of a problem subject to relaxed constraints necessarily includes the optimum of the original problem. Because the approximation is based on upper and lower envelopes, instead of a best-fit approximation, the feasible set in the new problem includes at least the feasible set of the original problem. □

The method provides valid bounds independent of how the pieces are defined (i.e., how the approximation breaks regions of the nonlinear function into piece-wise linear spaces). It also does not change the nature of equilibria, as no approximation of the underlying functions is involved (e.g., as opposed to log-linearizing equilibrium constraints). Indeed, one could generate several instances in which the points at which the pieces are considered are randomly generated.[14] The joint set of counterfactual bounds obtained across alternative approximations would still provide valid conservative bounds to the counterfactual of interest.

Proposition 3. *Define an approximation* $\underline{\mathbf{G}}_r(\mathbf{x}, \mathbf{u}_r, \theta)$ *and* $\overline{\mathbf{G}}_r(\mathbf{x}, \mathbf{u}_r, \theta)$*, in which the breakpoints for the integer variables* $\mathbf{u}_r$ *have been defined randomly. Consider the counterfactual bounds to a particular outcome generated by this approximation, given by* $[\underline{W}_{ur}(\theta), \overline{W}_{ur}(\theta)]$*. Consider several such approximations,* $r = 1, \ldots, R$*. Valid counterfactual bounds are given by* $[\max\{\underline{W}_{ur}(\theta)\}_{r=1}^R, \min\{\overline{W}_{ur}(\theta)\}_{r=1}^R]$.

Proof. By Proposition 2, for each r, $\overline{W}_{ur}(\theta) \geq \overline{W}(\theta)$ and $\underline{W}_{ur}(\theta) \leq \underline{W}(\theta)$, for $r = 1, \ldots, R$. Therefore, $\max\{\underline{W}_{ur}(\theta)\}_{r=1}^R \leq \underline{W}(\theta)$ and $\overline{W}(\theta) \leq \min\{\overline{W}_{ur}(\theta)\}_{r=1}^R$. □

The main point of Proposition 3 is to emphasize that because the counterfactual bounds are always valid conservative bounds, alternative approximating strategies do not compromise the validity of the solution. Indeed, they can be used as additional refinements. The result highlights how finding conservative bounds to counterfactual outcomes can be a powerful device, as its conservative nature opens the door to approximation refinements that are simple to implement.

Proposition 3 is reassuring, as it implies that the bounds will be valid even if we take a limited number of points. This is important, as reducing the number of integer variables effectively reduces the dimensionality of the problem in practice. In fact, given the integer nature of the problem, reducing the number of pieces can more than linearly reduce the computational time. Whereas the bounds on our counterfactual of interest are still valid, they might be very wide and uninformative.

Fortunately, an iterative approach can be used to refine the approximation while holding the number of approximation points fixed. Using the same bounds approach, one can compute the smallest and largest value of a particular strategic variable that can be sustained in equilibrium, by setting $W(\mathbf{x}, \theta) = x_i$. Once bounds on a strategic variable have been found, one can use them to narrow the range of approximation. Indeed, the range over which equilibrium constraints are approximated can be narrowed down successively,

[14] Certainly, more efficient schemes could be explored.

Figure 2: Successive Approximation of Equilibrium Constraints

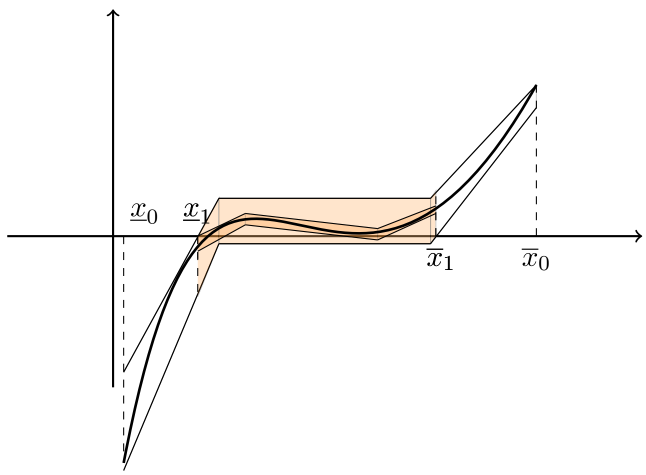

By narrowing down the relevant range of the action space, one can improve the approximation of equilibrium constraints, even when the number of piece-wise linear pieces is limited.

as suggested in Proposition 4. One key result is that such an iterative search does not invalidate the conservative nature of the counterfactual bounds. This contrasts with other iteration methods, which might converge to a particular region of the relevant strategy space.

Proposition 4. *Consider an iterative procedure in which the support of approximation for bounds on $\mathbf{G}(\mathbf{x};\theta)$ is given by $[\underline{\mathbf{x}}_k, \overline{\mathbf{x}}_k]$ at iteration k, and defined as $\underline{\mathbf{G}}_k(\mathbf{x},\mathbf{u},\theta)$ and $\overline{\mathbf{G}}_k(\mathbf{x},\mathbf{u},\theta)$. At iteration k, the counterfactual bounds on a given outcome are given by $[\underline{W}_{uk}(\theta), \overline{W}_{uk}(\theta)]$. At step $k+1$,*

1. *For $i = 1, ..., N$, where $N = \dim(\mathbf{x})$, define objective function $W = x_i$ and obtain bounds on variable $x_i \in [\underline{\omega}_{ik}, \overline{\omega}_{ik}]$.*
2. *For $i = 1, ..., N$, define new approximation bounds, $\underline{\mathbf{x}}_{i,k+1} = \underline{\omega}_{ik}, \overline{\mathbf{x}}_{i,k+1} = \overline{\omega}_{ik}$.*
3. *Compute new approximation bounds $\underline{\mathbf{G}}_{k+1}(\mathbf{x},\mathbf{u},\theta)$ and $\overline{\mathbf{G}}_{k+1}(\mathbf{x},\mathbf{u},\theta)$ under new support $[\underline{\mathbf{x}}_{k+1}, \overline{\mathbf{x}}_{k+1}]$.*
4. *Compute counterfactual bounds on the actual outcome of interest, $[\underline{W}_{u,k+1}(\theta), \overline{W}_{u,k+1}(\theta)]$.*

Iterated improvements using this approach provide valid counterfactual bounds. Furthermore, the bounds become (weakly) tighter as k increases, i.e., $\underline{W}_{uk}(\theta) \leq \underline{W}_{u,k+1}(\theta)$, and $\overline{W}_{uk}(\theta) \geq \overline{W}_{u,k+1}(\theta)$.

Proof. At each sub-iteration i at a given iteration k, one finds valid bounds to x_i. Therefore, restricting the approximating space does not exclude potential values of x_i that can be sustained in equilibrium. $\square$

Figure 2 provides a graphical intuition for this procedure in a single-dimensional space for ease of depiction.

Note that the procedure of maximizing and minimizing each strategic variable can also be helpful to ensure that constraints in the range of optimization are not binding, even if no iterations on the approximation range are performed. As an additional practical matter, note that depending on how costly step 3 is, one could do it in the inner loop i, instead of performing it only in the outer loop k. One could also pre-compute alternative approximations at different ranges, and iterate on those discrete intervals. The relative benefits of performing it for every ik, for every k, or in a more discrete manner might depend on the problem at hand.

2.3 Relaxation

The proposed optimization approach is amenable to alternative constraints, which can be useful when the equilibrium is difficult to solve. Indeed, the approximation method described above is already a particular case of a "relaxed" problem, in which the solution to a nonlinear game was transformed into a conservative set of piece-wise linear inequalities.[15] Similar to Proposition 2, any solutions to a relaxed problem will necessarily contain the original bounds of the problem.

Relaxing equilibrium conditions can be particularly useful in some applications due to computational complexity. In some applications, the original mixed-integer formulation might be too expensive. Therefore, a feasible solution to the problem might be complicated to find in a reasonable amount of time. Indeed, mixed-integer programs are NP-hard, making them increasingly burdensome as the number of variables grows. One kind of relaxed problem that is often considered in the optimization literature is one in which integer variables are transformed into continuous ones. This relaxed problem will still provide valid conservative bounds. How informative the bounds are will typically depend on the details of a given application. The advantage of transforming the problem into a linear program is that its computational complexity is decreased.

Observation 2. *Consider a relaxed integer program where $\mathbf{u} \in [0, 1]$. The relaxed problem provides valid conservative bounds to counterfactual outcomes. It has the advantage of transforming the problem from NP-hard to a linear program, at the expense of being less informative.*

The relaxation of the program can be especially beneficial in conjunction with Proposition 4. Consider a situation in which one tries to narrow down the support of a particular strategic variable, e.g., in a quantity game, the quantity produced by one of the firms. When searching for the maximum and minimum quantity for a given firm, one can use integer variables to approximate closely the first-order conditions of this particular firm, while relaxing the constraints that affect the other firms. This will ensure that the integer variables do not make the problem intractable, while still providing some information that can help tighten the bounds. We provide more details on how this insight can be used in the applications below.

2.4 Other extensions

Partial Equilibrium Characterization. For some models, we might be unable to fully characterize necessary and sufficient conditions for a game. This is quite common in the auction literature, where necessary

[15] In integer programming, the "relaxed" term is often used to refer to a linearized version of the problem, which convexifies integer variables. We use this term broadly, to define any changes in the constraints that make the problem more lax, i.e., the constraints become easier to be satisfied.

local conditions are easier to characterize than full equilibrium conditions. To the extent that necessary conditions for optimality are informative, they can be used to bound counterfactual outcomes. This is analogous to a relaxed problem: the bounds will still be valid, albeit larger.[16] Consider a game in which only necessary conditions can be characterized, or that necessary and sufficient conditions are known but too complex to describe. Necessary conditions to a game provide valid conservative bounds to counterfactual outcomes at the expense of being less informative. We show how to use this method in multi-unit auctions in Section 5.

Another class of problems that is easy to fit in this framework is situations in which players are not exactly optimizing, e.g., an epsilon equilibrium. Optimization errors or other perturbations to the equilibrium conditions can be easily incorporated by adding an epsilon term to both the upper and lower envelopes that define the equilibrium constraints.

Equilibrium Restrictions. Depending on the application, one might want to apply further restrictions to the equilibrium. For example, in particular economic applications, one might want to focus on equilibria that satisfy specific properties. Or maybe one would like to compare the counterfactual bounds to those that would arise in the presence of an explicit equilibrium-selection rule. Another natural restriction to add is second-order conditions, which rule out saddle points. In such cases, one can include additional constraints to the problem instead of relaxing it. Incorporating them into the framework is straightforward as long as those constraints can be bounded with piece-wise linear approximations and included in G. We highlight this application of the methodology in Section 4 by exemplifying how to use explicit refinements in Markov games.

Lack of Existence. Mixed-integer linear solvers try to efficiently assess all potential combinations in the program. Suppose the solver completes an exhaustive search and fails to find a set of equilibrium outcomes that are consistent with the relaxed equilibrium constraints. In that case, it can provide numerical proof that an equilibrium does not exist under such a set of constraints. Because the set of constraints is easier to satisfy than those of the original problem, it implies that it does not have an equilibrium. This highlights one of the advantages of the proposed methodology, which has the potential of being exhaustive. If the set of constraints were nonlinear, then we could only say that the solver did not find to manage a candidate equilibrium that satisfied the constraints. Still, we could not assert that no such candidate equilibrium existed. We also illustrate this application in 4.

Accounting for Parameter Uncertainty. Parameter uncertainty can also be easily incorporated in this framework, i.e. uncertainty in θ. In theory, one could solve for equilibrium bounds over several parameter combinations, bootstrapping from the estimated distribution of parameters. However, this can be very expensive. Fortunately, the framework allows one to compute conservative counterfactual bounds that account for parameter uncertainty. In particular, one can construct envelopes to equilibrium conditions that allow the parameters to take any value in the confidence interval or the identified set. Define Σ as the set in which

[16]The informativeness of these bounds depends on the particular application, as explored below.

parameter estimates are considered. Then, one can replace the envelopes to constraints with,

$$\underline{\mathbf{G}}(\mathbf{x},\mathbf{u},\Theta) = \min_{\theta\in\Theta}\ \underline{\mathbf{G}}(\mathbf{x},\mathbf{u},\theta),$$

$$\overline{\mathbf{G}}(\mathbf{x},\mathbf{u},\Theta) = \max_{\theta\in\Theta}\ \overline{\mathbf{G}}(\mathbf{x},\mathbf{u},\theta).$$

The counterfactual bounds obtained using this approach tend to be more conservative than a bootstrap approach in point-identified settings, as they ignore potential correlations between parameter estimates. Furthermore, they do not impose that the parameters take upon a particular value at each equilibrium equation simultaneously, but relax all equilibrium constraints to be consistent with the range of parameter estimates. In particular, one equilibrium constraint could be satisfied on the upper end of the confidence interval, whereas another could be satisfied at the lower end. However, the approach can be computationally much faster, as one only needs to compute the equilibrium of the game once, and it is still conservatively valid.

2.5 Discussion

The proposed approach presents some strengths and weaknesses that are important to discuss. The main strength of the methodology is that it is conservatively valid, i.e., counterfactual bounds always include the set of outcomes that can be sustained in equilibrium, although potentially more. This result is achieved thanks to the non-iteration and piece-wise linear approximation principles, that guarantee that the solution can be exhaustively checked for, as explained in Propositions 1 and 2. The approach also enables to search only for bounds, instead of characterizing all equilibria, which can have computational advantages.

The method relies heavily on approximation techniques. Naturally, the tightness of the counterfactual bounds will be as good as the approximation to the underlying game. Proposition 3 emphasizes that the particular choice of approximation does not invalidate the bounds, although it can affect their precision. It also opens the door to other approximation selection approaches. Proposition 4 establishes that the approximation range can be narrowed down iteratively while still ensuring the validity of the bounds.

Additionally, the method can work even when equilibrium conditions might be hard to fully characterize, under the principle of relaxation, as explained in Observation 2. A particular case of relaxation is when integer variables used for approximation are converted into continuous variables, substantially decreasing the computational burden of the problem, as explained in Observation 2.

The weaknesses of the method are related to its computational burden. Settings with highly nonlinear functions and many interactions will be hard to approximate with a mixed integer program. Whereas it is conceptually possible, the computational burden will increase exponentially. Observation 1 points out that, even if such limitations become binding, there is a practical way to account for the computational optimality error of the solution. However, it requires that at least one feasible combination is found.

To reduce the dimensionality of the problem, rougher approximation bounds can also be considered, although the informativeness of the counterfactual bounds may be compromised. Proposition 4 might help improve some of the precision of the bounds, but it may come at a substantial computational expense, as it requires iteration.

In sum, the practicality of the method is necessarily an empirical question, depending on the characteristics of the given application at hand. To gain intuition on the methodology and assess its performance in practice, we present three examples from the literature in which limitations to counterfactual simulations have been previously identified. In Section 3, we discuss two examples of static pricing games under discrete choice demand systems: discrete types mixed logit and conditional heteroskedasticity logit. Section 4 analyzes a dynamic game. Finally, we consider a multi-unit auction in Section 5.

3 Oligopolistic competition under discrete choice demand

Even though static price competition in differentiated product markets is one of the workhorse models of modern industrial organization, there is no general proof of existence or uniqueness under flexible demand systems. Our framework helps characterize equilibrium outcomes in such contexts.

Indeed, an extensive body of work has dealt with the existence and uniqueness of equilibrium prices in Nash-Bertrand competition with differentiated products. Regarding existence, Caplin and Nalebuff (1991)'s result establishes conditions over consumer preference that ensure that the single-product profit function is quasi-concave. Milgrom and Roberts (1990) shows existence by determining that many games with these characteristics are supermodular. When demand is logit or CES, they show that the game played by single-product firms has a unique equilibrium established using dominant diagonal arguments.[17] Mizuno (2003) extends Caplin and Nalebuff (1991) results for single-product firms, showing uniqueness for more flexible demand systems such as the nested logit. In the multiproduct case, profits need not be quasi-concave (Hanson and Martin, 1996) nor supermodular -or log-supermodular- (Nocke and Schutz, 2018), which limits the applicability of Kakutani and Tarksi's fixed point theorems. For these cases, Konovalov and Sándor (2010) applied Kellogg (1976)'s result to show that the multiproduct logit and CES models have a unique equilibrium. Similarly, Nocke and Schutz (2018) uses tools from aggregative games to establish the existence and uniqueness of a class of games that includes the multiproduct nested logit.

However, there are no general proofs of uniqueness -or existence- for more general specifications as the mixed logit when price sensitivity varies across consumers,[18] and it is easy to construct examples of multiplicity in simple models such as the logit with conditional heteroskedasticity, even for single-product firms (Echenique and Komunjer, 2007).

Although multiple equilibria can pose issues for estimation, Dubé et al. (2012) shows how to use an MPEC approach to estimate the likelihood function while remaining agnostic about which equilibrium is being played in the data. Nevertheless, counterfactual simulations are subject to multiplicity concerns, even conditional on having consistent estimates of the demand and supply parameters. In the following sections, we show how to address the latter challenge when the demand system creates multiplicity, taking the value of the fundamentals as given.

[17] Vives (1999) observes there are three basic approaches generally used to prove uniqueness in the literature: the "contraction" approach consists of showing the best replies are a contraction and is usually based on diagonal dominance arguments; the "univalence" approach finds conditions under which the system of FOC is one-to-one; while the "index theory" approach establishes conditions under which FOC have a unique solution.

[18] Aksoy-Pierson et al. (2013) provide conditions for equilibrium existence and uniqueness for a model with a discrete number of consumer types in a mixed multinomial logit. However, price sensitivity is constant across consumers in their model.

3.1 Discrete choice with consumer types

First, we focus on how to use our framework in discrete choice models with random coefficients, which have been widely used to perform counterfactual analysis (See Berry et al. (1995, 1999); Nevo (2000b); Petrin (2002); Leslie (2004), among others).

Consider a market with $j = 0, \ldots, J$ products, including the outside option indexed by zero. There are $n = 1, \ldots, N$ firms, each selling a particular set of products, $\mathcal{J}_n \subset \mathcal{J}$. The profit of firm n is given by,

$$\pi_n = M \sum_{j \in \mathcal{J}_n} (p_j - mc_j) s_j(\mathbf{p}),$$

where p_j is the product price, mc_j is the marginal cost to produce good j, s_j is the endogenous market share of product j, and M is the market size. Consumers choose one among the J products and the outside option. Consider a market approximated with $i = 1, \ldots, ns$ consumer types, with individual product-specific tastes μ_{ij}.

In order to solve the pricing equilibrium, we need to solve the equilibrium prices and shares in the market. Under the assumption that each consumer type can be represented with logit demand, the set of necessary conditions that characterize the equilibrium is given by,

$$\text{[Demand]} \qquad G_j^1 = s_j - \sum_{i=1}^{ns} \omega_i \frac{\exp(\delta_j + \mu_{ij})}{\sum_{m \in \mathcal{J}} \exp(\delta_m + \mu_{im})} = 0, \qquad \forall j = 0, \ldots, J$$

$$\text{[Supply]} \qquad G_j^2 = s_j + \sum_{r=1}^{J} (p_r - mc_r) \frac{\partial s_r}{\partial p_j} = 0, \qquad \forall j = 1, \ldots, J$$

where ω_i is a population weight for consumer type i, $\delta_j \equiv x_j \beta - \alpha p_j + \xi_j$ is a common valuation for product j across consumers, and $\mu_{ij} \equiv x_j \tilde{\beta}_i - \tilde{\alpha}_i p_j$ represents its idiosyncratic part.[19] The utility for product 0 (outside option) is typically normalized to zero. More compactly,

$$\mathbf{G}(\mathbf{s}^*, \mathbf{p}^*) = \begin{bmatrix} G_0^1(\mathbf{s}^*, \mathbf{p}^*) \\ G_1^1(\mathbf{s}^*, \mathbf{p}^*) \\ \vdots \\ G_J^2(\mathbf{s}^*, \mathbf{p}^*) \end{bmatrix} = 0.$$

In discrete choice models, researchers are often interested in the impact of a particular policy (e.g., a merger, a tax on imports, etc.) on consumer surplus. Imagine that we are interested in the following question. What is the maximum and minimum consumer surplus that can be attained in equilibrium, taking the estimated parameters as given and a particular policy change into account? Using the log-sum representation,

[19] See Nevo (2000a) for a detailed explanation of the demand side of the model.

we can solve for bounds on consumer surplus as,

$$\max_{\mathbf{s}^*,\mathbf{p}^*} \text{ and } \min_{\mathbf{s}^*,\mathbf{p}^*} \qquad \sum_i w_i \log\left(\sum_{j\in\mathcal{J}} \exp(\delta_j + \mu_{ij})\right)$$

$$\text{s.t.} \qquad \underline{\mathbf{G}}(\mathbf{s}^*,\mathbf{p}^*) \leq 0 \leq \overline{\mathbf{G}}(\mathbf{s}^*,\mathbf{p}^*).$$

where $\underline{\mathbf{G}}(\mathbf{s}^*,\mathbf{p}^*)$ and $\overline{\mathbf{G}}(\mathbf{s}^*,\mathbf{p}^*)$ are piece-wise linear relaxed conditions that conservatively bound $\mathbf{G}(\mathbf{s}^*,\mathbf{p}^*)$.

Note that the above formulation presumes that an equilibrium exists and that the necessary conditions are also sufficient. As we have discussed in the introduction of this section, an equilibrium may not exist for mixed logit systems. Thus, one should be careful when interpreting the counterfactual bounds, as it is under the assumption of equilibrium existence. Also, suppose one cannot find a set of outcomes satisfying the equilibrium bounds, conditional on the search being exhaustive. In that case, it necessarily implies that the equilibrium does *not* exist. Therefore, the method can provide numerical proof of *non-existence*, but not the opposite. Indeed, any bounds represent the minimum and maximum surplus that can be achieved when a relaxation of the necessary conditions of the problem is satisfied. These, however, may not fully determine an equilibrium of the game. In any case, if an equilibrium exists, its welfare lies within the bounds.

Extension to Discrete Choice with Endogenous Attributes Adding characteristic choices to the model can be achieved by introducing an additional vector of variables (characteristics of each car) and a set of first-order conditions for characteristic choices. Product characteristics enter into the consumer's utility as prices. The choice of characteristics has a similar impact to expected profits, and thus, its first-order condition can be added using similar techniques.

Consider a situation in which firms choose characteristic x_j, with the following normalized cost:

$$C(x_j) = \gamma_0 x_j s_j + \frac{\gamma_1}{2} x_j^2,$$

i.e., there is an increasing marginal cost of developing a given characteristic, which is independent of whether the product is successful or not, and also a cost per user of supplying it. One can add the additional first-order condition:

$$\text{[Supply } x\text{]} \qquad G_j^3 = \sum_{r=1}^{J} (p_r - mc_r - \gamma_0 x_r)\,\frac{\partial s_r}{\partial x_j} - \gamma_0 s_j - \gamma_1 x_j,$$

where the marginal cost includes now the costs of supplying characteristic x. Like the pricing case, the logit formulation gives explicit solutions to $\partial s_r/\partial x_j$. These constraints can be incorporated into the pricing and share constraints to compute equilibria with attribute choice. The method will be robust to multiple equilibria also along this dimension.

3.1.1 Envelopes on Equilibrium Constraints

How do we construct envelopes around the constraints in G? To build intuition, we start with the demand equations. It is helpful to note that, in the random-coefficients discrete choice framework, each consumer is solving a simple logit demand problem. One can define a set of auxiliary market shares s_{ij}, which represent the ex-ante probability of a given consumer i buying product j.

With the mean utility of the outside good normalized to zero, the following set of equations characterize the expected choice of each consumer type,

$$\begin{aligned}
0 &= \sum_j s_{ij} - 1, && \forall i = 1, \ldots, ns, \\
0 &= \ln(s_{ij}) - \ln(s_{i0}) - x_j\beta + \alpha\overline{p}_j - \xi_j - \mu_{ij}, && \forall i = 1, \ldots, ns, \forall j = 1, \ldots, J.
\end{aligned}$$

Constraints are already very linear. It is sufficient to bound the log function between 0 and 1 with a piece-wise linear function to obtain piece-wise linear bounds to the above constraints.[20] For each consumer i, the non-linear terms can be substituted with auxiliary variables γ_{ij}, and converted into piece-wise linear constraints as,

$$\begin{aligned}
0 &= \gamma_{ij} - \gamma_{i0} - x_j\beta + \alpha\overline{p}_j - \xi_j - \mu_{ij}, && \forall i = 1, \ldots, ns, \forall j = 1, \ldots, J, \\
\underline{\gamma}(s_{ij}) &\leq \gamma_{ij} \ \leq \overline{\gamma}(s_{ij}), && \forall i = 1, \ldots, ns, \forall j = 1, \ldots, J,
\end{aligned}$$

where $\overline{\gamma}$ and $\underline{\gamma}$ are a piece-wise upper and lower bound to the log function. With this reformulation, a set of piece-wise linear constraints defines the demand-side relaxed equilibrium constraints.

Additionally, to obtain aggregate shares, we add the linear constraints,

$$0 = s_j - \sum_i \omega_i s_{ij}, \quad \forall j = 0, \ldots, J.$$

We can use a similar technique to bound the supply-side equations as in the demand-side case. The key non-linear term in the first-order condition of the firms is the partial effect of prices on shares of products owned by the same firm. For a given product j, using the logit results,

$$\frac{\partial s_r}{\partial p_j} = \begin{cases} \sum_i \omega_i \alpha_i s_{ij}(s_{ij} - 1) & \text{if } r = j; \\ \sum_i \omega_i \alpha_i s_{ij} s_{ir} & \text{if } r \neq j \text{, and } j \text{ and } r \text{ belong to the same firm;} \\ 0 & \text{otherwise,} \end{cases}$$

with $\alpha_i = \alpha + \tilde{\alpha}_i$. These terms are clearly non-linear. Furthermore, they interact with product markups $p_r - mc_r$ in the first order conditions, so we need to approximate them.

To explain how to bound these terms, we focus on the own elasticity term; that is, we need to bound the

[20] Alternatively, one can define the log of the shares as a variable on its own, and bound the exponent of this variable to obtain the shares.

function $s_{ij}(1-s_{ij})$,

$$\underline{\zeta}(s_{ij}) \leq s_{ij}(1-s_{ij}) \leq \overline{\zeta}(s_{ij}),$$

which can be achieved with piece-wise bounds. Because shares were already used in the approximation of the log for the demand equations, this approximation does not increase the computational costs excessively.

However, $s_{ij}(1-s_{ij})$ also appears interacted with the markup $p_j - mc_j$. One way to approximate this interaction is to divide prices into bins. Within a given bin, the maximum price is given by the upper bound on the bin, similarly for the minimum price. Consider grid bin k. Then,

$$\underline{p}_k \underline{\zeta}(s_{ij}) \leq p_j s_{ij}(1-s_{ij}) \leq \overline{p}_k \overline{\zeta}(s_{ij}), \quad \text{if } p_j \in [\underline{p}_k, \overline{p}_k].$$

To improve the precision of the equilibrium bounds, one can increase the grid size of the markup approximation. Alternatively, one can use the results from Proposition 4 to narrow down the relevant range of markup levels iteratively. This option can be substantially less computationally intensive, as it reduces the dimensionality of the mixed integer program being solved at each iteration.[21]

Finally, we can add additional linear restrictions to complement the above restrictions and increase the algorithm's precision. For example, one can also bound the constraints with the following:

$$\begin{aligned}
&\underline{p}_k s_{ij}(1-\overline{s}_{ij}) \leq p_j s_{ij}(1-s_{ij}) \leq \overline{p}_k s_{ij}(1-\underline{s}_{ij}), \quad \text{if } p_j \in [\underline{p}_k, \overline{p}_k],\\
&\underline{p}_k \underline{s}_{ij}(1-s_{ij}) \leq p_j s_{ij}(1-s_{ij}) \leq \overline{p}_k \overline{s}_{ij}(1-s_{ij}), \quad \text{if } p_j \in [\underline{p}_k, \overline{p}_k],\\
&p_j \underline{s}_{ij}(1-\overline{s}_{ij}) \leq p_j s_{ij}(1-s_{ij}) \leq p_j \overline{s}_{ij}(1-\underline{s}_{ij}).
\end{aligned}$$

These additional constraints are linear and can help construct tighter envelopes to the equilibrium conditions.[22]

Given the envelopes around the equilibrium constraints, one can solve for upper and lower bounds to consumer surplus. Using the log-sum formulation and the fact that we have auxiliary variables γ_{ij} representing the log of the shares, we can get bounds on consumer surplus as,

$$\begin{aligned}
\max_{\mathbf{s}^*,\mathbf{p}^*} \text{ and } \min_{\mathbf{s}^*,\mathbf{p}^*} \quad & \sum_i -w_i \gamma_{i0} \\
\text{s.t.} \quad & \underline{\mathbf{G}}(\mathbf{s}^*, \mathbf{p}^*, \gamma^*, \mathbf{u}^*) \leq 0 \leq \overline{\mathbf{G}}(\mathbf{s}^*, \mathbf{p}^*, \gamma^*, \mathbf{u}^*),
\end{aligned} \tag{3}$$

where we use $\mathbf{u}^*$ to define the additional integer variables that create the different regions of the piece-

[21] In practice, We find that such an iterative procedure can be very successful at increasing the precision of the counterfactual bounds.

[22] Piece-wise linear bounds of quadratic terms can be handled automatically by several available optimizers. For instance, Gurobi can solve optimization programs up to global optimality, including cases with quadratic terms in equality restrictions.

wise linear constraints, see Appendix A.1 for implementation details.

3.1.2 A numerical example: Multi-Peak Consumers

We now illustrate how the method works through a numerical example. Consider the case of three consumer types and two products. Instead of having similar consumer types with varying sensitivities to product characteristics, consumers have substantial brand-specific tastes.[23] some consumers like product 1, consumers who like product 2, and a third set of more price-sensitive consumers, and we denote them as "shoppers". With these preferences, firms face a trade-off: exploit loyal consumers and set high prices or post low prices to capture the price-sensitive consumers.[24]

We decompose utility in just three terms:

$$u_{ij} = \delta_{ij} - \alpha_i p_j + \epsilon_{ij}$$

where ϵ_{ij} is a random variable with the type-I extreme value distribution iid across consumers and products, δ_{ij} are consumer-specific intercepts, and α_i is the consumer-specific price sensitivity coefficient. The utility of the outside good is normalized to 0. Table 1 presents the parameters used in the example.

Table 1: Parameter values for discrete types example.

	Type 1	Type 2	Type 3
δ_{i1}	8.00	11.00	4.60
δ_{i2}	11.00	8.00	4.60
α_i	-2.35	-2.35	[-4.00,-2.35]

Then, we solve the optimization program 3, using the parameter values in Table 1, to compute welfare bounds for several price sensitivities of the shoppers. To solve this optimization program we must compute piece-wise envelopes on the non-linear equilibrium constraints, i.e. $\underline{\mathbf{G}}(\mathbf{s}^*, \mathbf{p}^*, \gamma^*, \mathbf{u}^*), \overline{\mathbf{G}}(\mathbf{s}^*, \mathbf{p}^*, \gamma^*, \mathbf{u}^*)$.[25] To compare the method with more traditional approaches, we also solve the equilibrium by an iterative approach and back out consumer surplus from equilibrium prices. In Figure B.3, we also solve the non-linear MPEC to compare our results.

[23] A leading example of this type of demand is a model with structural state dependence. See, for example, Dube et al. (2010).

[24] This is the problem that much of the theoretical literature faces to prove the existence of an equilibrium. For instance, Caplin and Nalebuff (1991) noted: "*Without any restrictions on market demand, it may be that two extreme strategies, either charging a high price to a select group of customers (for whom the product is well positioned) or charging a low price to a mass market, both dominate the strategy setting an intermediate price. It is this issue which has been a major stumbling block in the study of existence*".

[25] One can approximate the solution to the demand system by computing piece-wise linear bounds to the logarithmic function between $(0, 1)$. Given the nature of the log function, it is helpful to concentrate the pieces at low values of s. Still, more sophisticated methods can be used to choose the location of the pieces, such as optimal knot placement routines developed in the splines literature (de Boor, 2001). Note that splines routines are geared at finding optimal knots for an approximation, not upper and lower envelopes, and thus, such "optimal" knots will not be optimal in the same exact sense. In practice, we iterate over the relevant range of shares to improve the approximation.

Recall that, in this context, we can neither theoretically guarantee equilibrium prices exist nor that constraints $G(s^*, p^*) = 0$ are sufficient to define the equilibrium. Thus, both welfare bounds and iterative solutions require careful analysis. Indeed, Figure B.1 illustrates the profit function is not quasi-concave for some shoppers' price sensitivity values, which is the mathematical characterization of firms' trade-off between pricing to a niche or the mass market. Moreover, it implies that prices satisfying FOC are not necessarily an equilibrium of the game. Actually, FOC can be satisfied even though firms are pricing at a relative minimum. This limitation is common to the iterative approach, which also searches for prices that solve FOC. However, our methodology allows us to add additional restrictions to Equation 3 to ensure second-order conditions are satisfied, ruling out the previous case. Nevertheless, even when firms are pricing at a relative maximum, it may not be the best response to rival strategies. There is no simple way to circumvent this challenge within our framework. Hence, we cannot rule out that the welfare bounds found are determined by outcomes (prices and shares) that are not an equilibrium of the game. In any case, bounds are always conservative. They always bound the minimum and maximum welfare that can be achieved in any equilibrium if such an equilibrium exists. In this simple case, we can compute all *actual* equilibria of the game, intersecting firms' best responses (see Figure B.2). Hence, we compare the solution found by each approach with the actual equilibrium of the game.

Figure 3 depicts the bounds on consumer surplus (CS) obtained by solving Equation 3, the true equilibrium CS, and the solution found using a traditional iterative approach. The blue dots in Figure 3 indicate that the multiplicity of equilibria depends on shoppers' price sensitivity. When shoppers are not very sensitive, the unique equilibrium sets low prices to target all consumers in the market, implying a high CS equilibrium. As shoppers get more price sensitive, the incentives to fight over them decrease. Indeed, there is a range of price sensitivity parameters where both strategies are an equilibrium, and both high and low CS can arise (indicated with grey dotted vertical lines in Figure 3)). Past this range, the unique equilibrium is to give up on shoppers and focus on "captive" consumers.

Figure 3: Bounds on equilibrium consumer surplus with heterogeneous consumer's type.

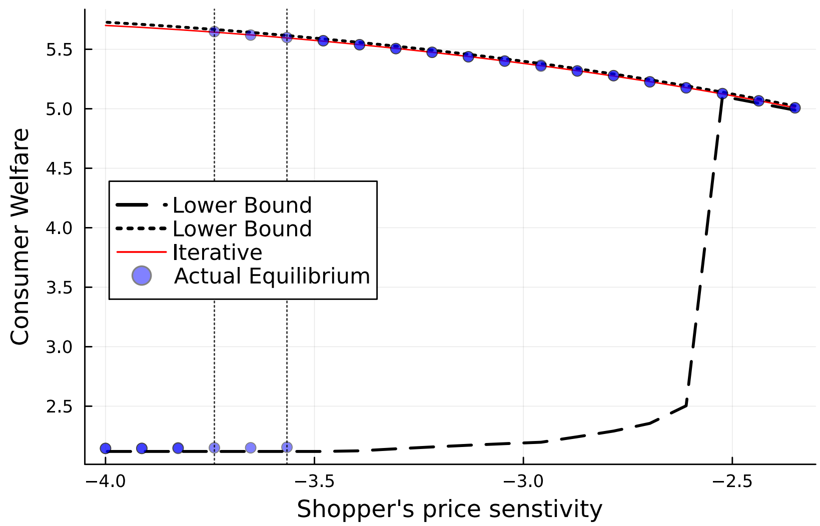

Notes: Figure B.3 show the results for a non-linear solution of Equation 3, when we substitute the piece-wise linear restrictions $\underline{\mathbf{G}}(\mathbf{s}^*, \mathbf{p}^*, \gamma^*, \mathbf{u}^*) \leq 0 \leq \overline{\mathbf{G}}(\mathbf{s}^*, \mathbf{p}^*, \gamma^*, \mathbf{u}^*)$ by the true non-linear necessary conditions $\mathbf{G}(\mathbf{s}^*, \mathbf{p}^*) = 0$.

Although CS bounds always contain the true equilibria of the game, sometimes we compute wide bounds despite the equilibrium being unique. This indicates that only a subset of vectors satisfying first and second-order conditions is an equilibrium of the game. Wide bounds can still be used to alert of this possibility.

Second, observe that the iterative approach ”lands” on the high welfare equilibrium (this is dependent on initial conditions). In particular, the iterative approach continues to reach this solution even when it no longer constitutes an equilibrium of the game, i.e. for large price sensitivities. As we mentioned, the iterative solution can even converge to prices where some firms are pricing at a relative minimum, which is ruled out within our framework. Conservative bounds can again warn of this possibility.

Third, welfare bounds are tight only when a unique price vector satisfies FOC and SOC. In this case, the iterative approach lies between the computed bounds, which can be used to prove uniqueness numerically. This also puts Proposition 2 in perspective: although bounds are conservative with respect to sharp bounds, they can be set arbitrarily close to the sharp bounds by bounding the error of the piece-wise linearization.

Finally, Panels A1 and B1 in Table 2 show numerical values of the CS bounds -and its corresponding equilibrium shares and prices- for two values of shoppers' price sensitivity: when equilibrium is unique $\hat{\alpha} = -2.524$, and under multiplicity $\hat{\alpha} = -3.653$. Panel A1 illustrates our previous observation: when a unique vector of prices satisfies FOC and SOC, the distance between the upper and lower bound can be made arbitrarily narrow. Hence, the presence of wide bounds, as in the case of panel B1, may alert the possibility of multiple equilibria, although it does not prove it.

On the other hand, panels A2 and B2 present bounds on market shares and markups. We compute bounds on each market outcome one at a time, with the only constraint that market shares and markups for other firms are consistent with the relaxed equilibrium conditions. Naturally, not all bounds need to be sustained at the same time. For instance, the sum of market shares' lower bound does not need to add up to

1. For this reason, it is still useful to bound the economic outcome of interest, such as CS (shown in Panel B1 and A1 in Table 2, and Figure 3).

In addition, when we bound CS directly, we do not expect equilibrium shares and prices in high and low CS scenarios to hold any particular relation, nor do prices or shares in the iterated approach need to lay within the equilibrium outcomes of these variables in each CS scenario. However, any equilibrium share or markup (either in the proposed approach or iterative approaches) must be between the bounds established in Panels A2 and B2, as is the case in Table 2.

3.2 Discrete choice with conditional heteroskedasticty

The previous sections showed how to relax the equilibrium conditions of an oligopoly pricing game by constructing piece-wise envelopes of the logarithm (or exponential) function and bounding the quadratic interactions between variables. It was emphasized, however, that the "amount" of relaxation needed was small because the original problem was almost linear. Indeed, the introduction of multiple types of consumers did not add any non-linearity since equilibrium conditions are obtained by adding restrictions for each type.

Although we have shown that the proposed method works well in such a case, one may be concerned about its computational tractability once we move toward less linear environments. Next, we illustrate how the method works when multiple equilibria arise from conditional heteroskedasticity in the logit error, as shown in Echenique and Komunjer (2007).

In this example, two identical firms produce a single product and no outside option. Consumers' utility is defined as

$$u_{ij} = -\alpha p_j + \nu_{ij}$$

Hence, all consumers have the same price sensitivity. Furthermore, the error variance depends on the price level. Concretely, $\nu_{ij} = g(p_j p_{-j})\epsilon_{ij}$, where ϵ_{ij} are Gumbel random variables, iid across firms and consumers, with unit variance. In our example, $g(x)$ is a third-degree polynomial. That is, the overall price level of the market affects the degree of horizontal differentiation across consumers. On the firm side, we assume constant marginal costs, which are equal across firms: $ln\ c_j = \gamma + ln\ \tau_{tax}$, where τ_{tax} is a tax whose value affects the nature of the equilibrium outcomes.[26]

Under these conditions, the market share of product j, which also characterizes the first set of necessary conditions, $H^1_j(p,s)$, is

$$\text{[Demand]} \qquad H^1_j(p,s) = s_j(p_j, p_{-j}) - \frac{1}{1 + exp(\frac{\alpha(p_j - p_{-j})}{\sigma g(p_j p_{-j})})} = 0, \quad j \in \{1, 2\}$$

where $\sigma = \sqrt{6}/\pi$.

Observe that the $g(p_j p_{-j})$ term makes this equilibrium condition "less" linear than in our previous

[26] Differently from our previous example, in this case, we can show the existence of equilibrium using Caplin and Nalebuff (1991) result, since the density function of individual preferences is log-concave.

Table 2: Bounds on Consumer Surplus, Shares and Markups (Example 1)

Panel A1: Unique Equilibrium, $\hat{\alpha}$ = -2.524

	Lower	Iter	Upper	Iter−Low	Upper−Iter
Consumer Surplus	5.1245	5.1248	5.1285	2.8e-04	3.7e-03
Eqm. Shares					
Outside Option	0.0913	0.0913	0.0911		
Product 1	0.4543	0.4543	0.4544		
Product 2	0.4543	0.4543	0.4544		
Eqm. Markups					
Product 1	1.7077	1.7076	1.7062		
Product 2	1.7077	1.7076	1.7062		

Panel A2: Unique Equilibrium, Bounds on Prices and Shares

	Lower	Iter	Upper	Iter−Low	Upper−Iter
Shares					
Outside Option	0.0911	0.0913	0.0913	2.1e-04	3.1e-05
Product 1	0.4542	0.4543	0.4545	1.0e-04	1.9e-04
Product 2	0.4542	0.4543	0.4545	1.0e-04	1.9e-04
Markups					
Product 1	1.7062	1.7076	1.7077	1.4e-03	1.2e-04
Product 2	1.7062	1.7076	1.7078	1.4e-03	1.8e-04

Panel B1: Multiple Equilibria, $\hat{\alpha}$ = -3.653

	Lower	Iter	Upper	Iter−Low	Upper−Iter
Consumer Surplus	2.1501	5.6207	5.6242	3.5e+00	3.5e-03
Eqm. Shares					
Outside Option	0.3595	0.1278	0.1274		
Product 1	0.3202	0.4361	0.4363		
Product 2	0.3202	0.4361	0.4363		
Eqm. Markups					
Product 1	3.3466	1.3183	1.3169		
Product 2	3.3466	1.3183	1.3169		

Panel B2: Multiple Equilibira, Bounds on Prices and Shares

	Lower	Iter	Upper	Iter−Low	Upper−Iter
Shares					
Outside Option	0.1274	0.1278	0.3596	3.9e-04	2.3e-01
Product 1	0.3201	0.4361	0.4365	1.2e-01	4.3e-04
Product 2	0.3201	0.4361	0.4365	1.2e-01	4.3e-04
Markups					
Product 1	1.3167	1.3183	3.3466	1.6e-03	2.0e+00
Product 2	1.3167	1.3183	3.3466	1.6e-03	2.0e+00

Notes: Results from Example 1. Panel A features a parameterization with two products and a unique equilibrium. Panel B features a parameterization with two products and three consumer types (loyal to product 1, loyal to product 2, and non-loyal), which features multiple equilibria.

examples. Thus, constructing piece-wise linear bounds on the exponential function is insufficient to bound equilibrium conditions. Nevertheless, we can still follow a similar approach, defining auxiliary variables and constructing piece-wise envelopes of the exponential function for shares and piece-wise envelopes of a third-degree polynomial for the conditional variance function and of a second-degree polynomial for its derivative. Even though interactions between variables remain after such relaxation, we can further relax quadratic interactions exactly as we did in the previous examples.[27]

On the supply side, non-linearities are also more prevalent. Now, the partial effect of firms' prices on shares is

$$\frac{\partial s_j}{\partial p_j} = -\alpha \frac{s_j(1-s_j)}{\sigma g(p_1 p_2)^2} \left\{ g(p_1 p_2) + (p_{-j} - p_j) \frac{\partial g(x)}{\partial x}|_{x=p_1 p_2 p_{-j}} \right\}, \quad j = \{1, 2\}$$

Again, we can relax this constraint by including appropriate auxiliary variables and constructing piece-wise linear bounds on polynomial and exponential functions. Then, the FOC of the firms are

$$\text{[Supply]} \qquad H_j^2(p, s; \delta) = s_j + \frac{\partial s_j}{\partial p_j}(p_j - c_j(\tau_{tax})) = 0, \quad j = \{1, 2\}$$

As before, the set of necessary conditions characterizing the equilibrium can be written succinctly as

$$\mathbf{H}(\mathbf{s}^*, \mathbf{p}^*; \delta) = \begin{bmatrix} H_1^1(\mathbf{s}^*, \mathbf{p}^*) \\ H_2^1(\mathbf{s}^*, \mathbf{p}^*) \\ H_1^2(\mathbf{s}^*, \mathbf{p}^*; \delta) \\ H_2^2(\mathbf{s}^*, \mathbf{p}^*; \delta) \end{bmatrix} = 0.$$

3.2.1 A numerical example: implementation in a highly non-linear environment

Appendix A.2 shows how to construct envelopes around these equilibrium constraints and describes a specific implementation using available optimization software. Given these envelopes, which depend on prices, shares, a larger set of auxiliary variables, Σ, and model parameters (in particular τ_{tax}), we can get bounds on any equilibrium outcome. In the following example, we find the maximum and minimum price that can be sustained in equilibrium for several tax levels:

$$\begin{aligned} \bar{p_1}(\tau_{tax}) \text{ and } \underline{p_1}(\tau_{tax}) = \max_{\mathbf{s}^*, \mathbf{p}^*} \text{ and } \min_{\mathbf{s}^*, \mathbf{p}^*} \qquad & p_1 \\ \text{s.t.} \qquad & \underline{\mathbf{H}}(\mathbf{s}^*, \mathbf{p}^*, \mathbf{\Sigma}^*; \tau_{tax}) \leq 0 \leq \overline{\mathbf{H}}(\mathbf{s}^*, \mathbf{p}^*, \mathbf{\Sigma}^*; \tau_{tax}) \end{aligned} \tag{4}$$

We solve Equation 4 for the parameters shown in Table 3

Figure 4 presents the bounds $\bar{p_1}(\tau_{tax})$ and $\underline{p_1}(\tau_{tax})$ for several values of τ_{tax}. When taxes are low, the upper and lower bounds are tight, suggesting a unique equilibrium exists. As taxes increase, multiple

[27]Furthermore, available optimizers solve mixed-integer problems with non-convex quadratic constraints to global optimality. Hence, quadratic interactions can be handled internally by available software.

Table 3: Parameter values for conditional heteroskedastic logit example.

	Parameters
α	0.39
$g(x)$	$0.1520 + -0.0030x + 0.0089(x^2) + -0.0005(x^3)$
γ	-0.10
$ln(\tau_{tax})$	[0.00,0.20]

Notes: The coefficients of the polynomial are chosen to approximate the conditional variance function in Echenique and Komunjer (2007): $g(x) = \phi + exp(-\tau/x)$, where $\rho = 0.166$ and $\tau = 10$.

Figure 4: Bounds on equilibrium prices with conditional heteroskedastic logit demand.

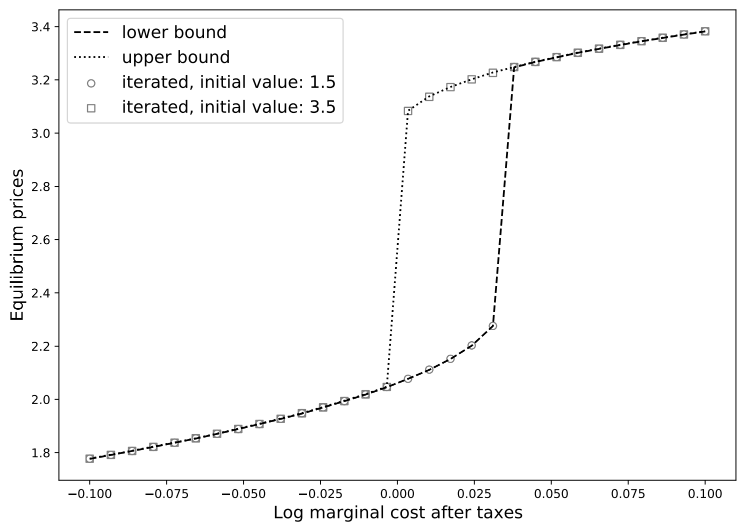

equilibria arise. However, uniqueness is reestablished in a higher price equilibrium when costs increase further. This example highlights the proposed method's utility in analyzing counterfactuals under multiple equilibria. Suppose a market is operating at a low cost such that the equilibrium is unique but close to the multiplicity threshold. Then, a regulator considers imposing a tax on firms, shifting overall costs slightly above the threshold. In such a case, the tax could have a small effect on prices or a large and discrete effect, if the equilibrium selection mechanism pushes firms toward the high price equilibrium.

Regarding the quality of the equilibrium relaxation, Figure 4 suggests the information loss is small since we obtain tight bounds in the region under uniqueness. In this case, an exhaustive exploration of equilibria modifying the initial conditions within iterative algorithms presents a similar picture to our bounds.[28]

Next, we show how to use our methodology to explore intermediate equilibria, i.e. equilibrium prices lying strictly within computed bounds. In principle, we could do this by either imposing additional restric-

[28]In addition, Figure B.4 presents the outcome of the non-linear MPEC problem, where the constraint is given by $\mathbf{H}(\mathbf{s}^*, \mathbf{p}^*; \delta) = 0$, instead of the relaxed equilibrium conditions. We can see that non-linear solvers generally cannot discriminate between low-price and high-price equilibrium when multiple equilibria are present.

Figure 5: Bounds on equilibrium prices within a conditional heteroskedastic logit demand.

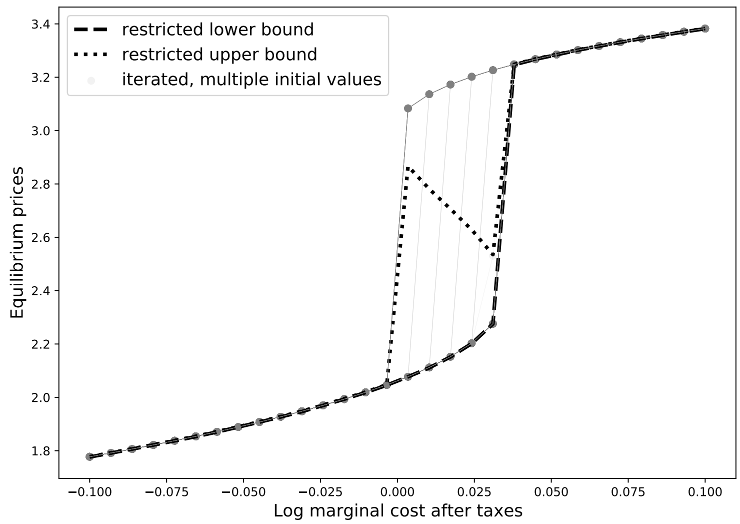

Notes: Bounds are constructed imposing further restrictions to Equation 4. In particular, we impose that the upper bound must be strictly lower than the unrestricted bound presented in Figure 4.

tions on optimization program Equation 4, narrowing down the search with non-linear solvers, or changing initial conditions in usual iterative procedures. Figure 5 presents the result from this exercise. Here, we restrict bounds in the multiple equilibria region to be strictly lower than the upper bound in the unrestricted case. Moreover, we solve the iterative procedure for 100 initial values between 1.5 and 3.5. Our bounds suggest the existence of an intermediate price equilibrium within the multiplicity region, which is decreasing costs.[29] Interestingly, the iterative procedures cannot find this equilibrium despite the large number of initial values chosen. It is well known that iterative procedures may be unable to find unstable equilibria.[30]

3.3 Discussion

The above examples illustrate how to robustly compute counterfactual outcomes of oligopoly pricing games with discrete choice demands. Given the potential multiplicity of equilibria in these models, we provide a valuable tool to perform a conservative analysis of policy impacts. This approach can be less appealing in settings with many products or consumer types, as the computational requirements can grow substantially. A scaled-down version of a model can still be useful in situations where iteration methods prove to be sensitive to starting values or fail to converge. It can also provide ballpark starting values for such iteration methods. Furthermore, recall that the method might overstate the size of the bounds but always contains all valid equilibria within the bounds. Failing to find an equilibrium using a coarse approximation can also help detect issues with the model parameterization (e.g., the lack of existence of an equilibrium in pure strategies).

[29]This equilibrium was also highlighted by Echenique and Komunjer (2007).
[30]Interestingly, non-linear methods do find this equilibrium -see Figure B.4.

4 Dynamic Games

A commonly used framework for industry dynamics with strategic interactions is based on a Markov assumption (Maskin and Tirole, 1988; Ericson and Pakes, 1995). This framework has been used to model various industries (Benkard, 2004; Jofre-Bonet and Pesendorfer, 2003; Ryan, 2012). Whereas dynamic interactions are relevant in strategic settings, implementation of models of industry dynamics has been limited. Both the computational burden of computing Markov equilibria and our narrow capability to characterize all such equilibria exhaustively have contributed to its scarce application.

To address the first problem, researchers have proposed relaxing the equilibrium concept (Weintraub et al., 2008; Fershtman and Pakes, 2012; Abbring et al., 2016) or relying on approximating techniques (Farias et al., 2012). To solve the second issue, some authors have proposed homotopy methods. Although these methods constitute a decisive step forward in characterizing multiplicity, they still present several limitations. On the one hand, the classic homotopy method cannot find all equilibria. On the other hand, the all-solution homotopy can only be applied to games whose equilibrium conditions are characterized by a few dozen polynomial functions (Judd et al., 2012).

The tools presented in this paper aim to address the second issue, complementing homotopy methods. They can also be applied to alternative, less computationally demanding equilibrium concepts, such as oblivious or moment-based equilibria, to the extent that multiple equilibria are still a concern in such a setting.[31] Finally, we can use the approach to partially characterize equilibria in dynamic games in a less computationally intensive way, as shown below.

4.1 Model outline

Consider the dynamic game on learning-by-doing and organizational forgetting presented in Besanko et al. (2010), which features multiple equilibria.[32] The game has two firms $n \in \{1, 2\}$, each with a current state $e_n \in \{1, ..., M\}$. This state reflects firms' cost function: the higher the states, the lower the firm's marginal cost is. In particular, the cost function is given by

$$c(e_n) = \begin{cases} \kappa e_n^\eta & \text{if } 1 \leq e_n < m, \\ \kappa m^\eta & \text{if } m \leq e_n \leq M, \end{cases}$$

where $\kappa > 0$, $\eta = \log_2 \rho$, and $\rho \in (0, 1]$. The probability of forgetting is given by $\Delta(e_n) = 1 - (1 - \delta)^{e_n}$, with $\delta \in [0, 1]$.

At any point in time, the industry is characterized by the vector of firms' states $\mathbf{e} = (e_1, e_2) \in \{1, ..., M\}^2$. The evolution of the state is modeled as follows:

$$e'_n = e_n + q_n - f_n,$$

[31] The method can be used as long as there is a set of necessary and/or sufficient set of equilibrium conditions that can be expressed as a system of equations.

[32] This subsection borrows heavily on the original paper Besanko et al. (2010). The goal is to show how to use the proposed methodology in a game in which homotopy methods have already characterized several equilibria.

where $q_n \in \{0, 1\}$ indicates whether a firm makes a sale and learns, and $f_n \in \{0, 1\}$ tells whether a firm forgets know-how.

The probability of a successful sale is a function of firms' pricing strategies. In particular,

$$Pr(q_n = 1) = D_n(\mathbf{p}) = \frac{1}{1 + \exp\left(\frac{p_n - p_{-n}}{\sigma}\right)},$$

where p_n is firm n price and $\sigma > 0$. In such cases, firms learn and reduce their marginal cost, i.e. their state increases by one with $q_n = 1$. However, the firm can still forget, which is given by f_n. The probability of forgetting is defined as,

$$Pr(f_n = 1) = 1 - (1 - \delta)^{e_n},$$

in which case, the firm loses some of the accumulated stock and suffers an increase in marginal cost, i.e., a reduction in its state.

Static profits are given by the probability of making a successful sale times the profit margin at that particular state, i.e.,

$$\Pi_1(\mathbf{e}) = D_1^*(\mathbf{e})\Big(p_1(\mathbf{e}) - c(e_1)\Big).$$

We focus on symmetric Markov-perfect equilibria. Following Besanko et al. (2010), the equilibrium of the game satisfies the following Bellman and FOC conditions:

$$\begin{aligned}
\text{[Bellman]} \qquad & F^1_{\mathbf{e}}(\mathbf{V}^*, \mathbf{p}^*) = -V^*(\mathbf{e}) + D_1^*(\mathbf{e})\left(p_1^*(\mathbf{e}) - c(e_1)\right) \\
& \qquad + \beta \sum_{k=1}^{2} D_k^*(\mathbf{e}) \overline{V}_k^*(\mathbf{e}) \\
& \quad = 0, \\
\text{[FOC]} \qquad & F^2_{\mathbf{e}}(\mathbf{V}^*, \mathbf{p}^*) = \sigma - \left(1 - D_1^*(\mathbf{e})\right)\left(p_1^*(\mathbf{e}) - c(e_1)\right) - \beta \overline{V}_1^*(\mathbf{e}) \\
& \qquad + \beta \sum_{k=1}^{2} D_k^*(\mathbf{e}) \overline{V}_k^*(\mathbf{e}) \\
& \quad = 0,
\end{aligned}$$

where the asterisks denote an equilibrium and $\overline{V}_k^*$ is the expectation of the value function of the firm conditional on the buyer purchasing its own good ($k = 1$) or the one of the other firm ($k = 2$). More compactly,

$$\mathbf{F}(\mathbf{V}^*, \mathbf{p}^*) = \begin{bmatrix} F^1_{(1,1)}(\mathbf{V}^*, \mathbf{p}^*) \\ F^1_{(2,1)}(\mathbf{V}^*, \mathbf{p}^*) \\ \vdots \\ F^2_{(M,M)}(\mathbf{V}^*, \mathbf{p}^*) \end{bmatrix} = 0.$$

The first-order conditions define a unique optimal price, and the problem satisfies conditions for equilibrium existence (Doraszelski and Satterthwaite, 2010; Besanko et al., 2010).

Finding all solutions to this system of non-linear equations is complicated and requires sophisticated homotopy methods. As an alternative, with the proposed method, we find bounds to the solution of the system of equations using approximation techniques. As explained above, the system of non-linear equations is substituted by conservative piece-wise linear constraints.

One possible approach is to directly bound V^* and P^* for every individual state. In that way, we characterize the range of prices and value functions that can be obtained in any equilibrium. In such case, we would need to solve two optimization programs for every variable (V^*, P^*) and state (s, k),

$$\begin{aligned} \max_{\mathbf{V}^*,\mathbf{p}^*} \text{ and } \min_{\mathbf{V}^*,\mathbf{p}^*} \quad & V^*(s,k) \text{ or } P^*(s,k) \\ \text{s.t.} \quad & \underline{\mathbf{F}}(\mathbf{V}^*,\mathbf{p}^*) \leq 0 \leq \overline{\mathbf{F}}(\mathbf{V}^*,\mathbf{p}^*), \end{aligned} \tag{5}$$

where $\underline{\mathbf{F}}(\mathbf{V}^*, \mathbf{p}^*)$ and $\underline{\mathbf{F}}(\mathbf{V}^*, \mathbf{p}^*)$ are piece-wise linear relaxed conditions that conservatively bound the actual equilibrium conditions $\mathbf{F}(\mathbf{V}^*, \mathbf{p}^*)$.

The computational burden of this approach is high since we have to solve a complicated system $M^2 \times 2 \times 2$ times. However, we can reduce it by decreasing the number of equilibrium conditions we include in any single optimization program, and solve the problem from states that can be easily bound, given previously found bounds, towards harder ones. This procedure represents a partial integration of our method with iterative solutions. That is, we cycle throughout the state space updating bounds on policies and values, using simple mixed integer programming programs that incorporate values and prices' bounds from previous iterations and the current iteration's previously visited states. See appendix A.3 for the detailed algorithm used to perform this procedure efficiently.

This approach constructs the lower and upper bound on prices and value functions that can be attained in any equilibria for every state. It does not mean that any two of such bounds must belong to the same underlying equilibrium. That is, every bound in a different state may be achieved under a different equilibrium. Alternatively, we can find bounds to particular outcomes of interest. For example, in the context of games with learning-and-forgetting, which has a race feature, it might be interesting to assess the minimum and maximum net present value at the beginning of the industry when both firms have high marginal costs and have not learned yet, i.e., when the state is equal to zero for both firms. In this case, we can solve two separate programs,

$$\begin{aligned} \max_{\mathbf{V}^*,\mathbf{p}^*} \text{ and } \min_{\mathbf{V}^*,\mathbf{p}^*} \quad & V^*(0,0) \\ \text{s.t.} \quad & \underline{\mathbf{F}}(\mathbf{V}^*,\mathbf{p}^*) \leq 0 \leq \overline{\mathbf{F}}(\mathbf{V}^*,\mathbf{p}^*), \end{aligned}$$

It is worth noting that even though we bound $V(0, 0)$, the program still solves for all variables, $\{\mathbf{V}^*, \mathbf{p}^*\}$. However, the method focuses on finding the upper and lower bound to the value function at that particular state while finding values for all the other variables consistent with the relaxed equilibrium constraints.

In the following sections, we show how to relax equilibrium conditions of Besanko et al. (2010)'s game and find bounds for the policies and values that can be achieved in any game's equilibrium. In the numerical examples, we fix most of the parameters in Besanko et al. (2010)'s model–see Table 4–and show how bounds on equilibrium outcomes vary as we change the discount factor and the forgetting parameter.

Table 4: Baseline parametrization for dynamic game computations.

Parameter	Description	Value
σ	Love for Variety	0.85
κ	Learning Curve - Level	1.0
ρ	Learning Curve - Steepness	10.0
M	States	8
m	Bottom learning curve	5

4.2 Envelopes on FOC

Similar to the discrete choice case, it is useful to show how the approximation technique works with the simplest condition, which in this case, is the first-order condition.

The FOCs at each state $\mathbf{e} = \{e_1, e_2\}$, are given by the non linear expressions,

$$\begin{aligned} 0 &= \frac{\sigma}{1 - D_1(p_1 - p_2)} - \left((p_1 - c_1) + \beta\left(\overline{V}_{11} - \overline{V}_{12}\right)\right), \\ 0 &= \frac{\sigma}{1 - D_2(p_2 - p_1)} - \left((p_2 - c_2) + \beta\left(\overline{V}_{21} - \overline{V}_{22}\right)\right). \end{aligned}$$

Most of the price and value function terms enter these conditions linearly. The only non-linear part is given by $p_1 - p_2$, which enters the demand function. We can transform these equations into piece-wise linear inequalities. Define the auxiliary variable $z \equiv p_1 - p_2$, $\overline{f}(z)$ as an upper bound to $\frac{\sigma}{1-D(z)}$, with $D(z) = \frac{1}{1+\exp\left(\frac{z}{\sigma}\right)}$, and $\underline{f}(z)$ as a lower bound to the same function. For a given state $\mathbf{e}$, the following system determines the solution to the game $\{p_1, p_2, z\}$,

$$\begin{aligned} \overline{f}(z) - (p_1 - c_1) &\geq 0, \\ \overline{f}(-z) - (p_2 - c_2) &\geq 0, \\ \underline{f}(z) - (p_1 - c_1) &\leq 0, \\ \underline{f}(-z) - (p_2 - c_2) &\leq 0, \\ p_1 - p_2 - z &= 0. \end{aligned}$$

Due to the approximation embedded in the system, the counterfactual bounds obtained with this method contain at least the actual solution. As the approximation becomes tighter, the bounds converge to the equilibrium of the game, which in this example, is unique. Exploring the size of the bounds as the number of approximation points increases can help get a sense of the accuracy of the bounds, absent any multiple equilibria.

4.2.1 Numerical example for the static case.

First, observe that when $\beta = 0$, we are in the static case, and solving for the FOC is equivalent to solving for the game's equilibrium. We solve for bounds at each state $\mathbf{e}$ by bounding the function $f(z) = \frac{\sigma}{1-D(z)}$ with a piece-wise linear approximation with J equal pieces between $\underline{z}$ and $\overline{z}$]. As $J \to \infty$, the upper and the lower bound converge to the unique equilibrium of the game. Figure B.6 shows the approximation to the function $\frac{\sigma}{1-D(z)}$ when $J = 5$ and $z \in [-2, 2]$. Given the well-behaved nature of the function, the approximation is quite accurate.

With such an approximation of f, one can define envelopes on the equilibrium conditions that are piece-wise linear and compute counterfactual bounds in a robust fashion. In this case, we compute bounds on equilibrium prices. Table 5 shows the method's accuracy for several node sizes and initial bounds. We also compare the counterfactual bounds to the actual unique equilibrium.

Table 5: Counterfactual bounds on prices for the static case

	$J = 5$ $z \in [-2, 2]$	$J = 5$ $z \in [-3, 3]$	$J = 10$ $z \in [-2, 2]$	$J = 10$ $z \in [-3, 3]$
Average bounds $[\underline{p}_1, \bar{p}_1]$	[9.71, 9.77]	[9.71, 9.77]	[9.71, 9.77]	[9.71, 9.77]
Iterated Prices p_1	9.72	9.72	9.72	9.72
Average $(\bar{p}_1 - \underline{p}_1)$	0.07	0.07	0.07	0.07
Maximum $(\bar{p}_1 - \underline{p}_1)$	0.10	0.10	0.10	0.10
Average $(\bar{p}_1 - \underline{p}_1)/\underline{p}_1$	0.01	0.01	0.01	0.01
Maximum $(\bar{p}_1 - \underline{p}_1)/\underline{p}_1$	0.01	0.01	0.01	0.01

Notes: Parameters are as in Table 4, with $\beta = 0$. The table presents averages across M^2 states. Note that in the static case, it is easy to characterize the unique equilibrium prices at each state using an iterative approach, which is also presented in the table as a reference point.

Table 5 illustrates that the bounds provided are relatively tight, even with relatively coarse approximations. In all four cases considered, there is at most a 1% to 2% difference between the upper and the lower bound. Nevertheless, as we increase the number of nodes and shrink the support of the approximation, we also improve the performance of the equilibrium bounds.

4.3 Envelopes on Bellman

As we move to the dynamic case, the system of equations that defines the equilibrium also includes the Bellman equations. Making use of the fact that $D_1(z) + D_2(-z) = 1$, the Bellman conditions become, [33]

$$0 = \frac{V_1 - \beta \overline{V}_{12}}{D_1(z)} - \left((p_1 - c_1) + \beta\left(\overline{V}_{11} - \overline{V}_{12}\right)\right),$$
$$0 = \frac{V_2 - \beta \overline{V}_{22}}{D_2(-z)} - \left((p_2 - c_2) + \beta\left(\overline{V}_{21} - \overline{V}_{22}\right)\right).$$

The non-linear terms affect both prices and value functions. In particular, $D_n(z)$ interacts with the value functions, in the form $g(z, V) = \frac{V}{D_n(z)}$, where V might be the value function at the particular state, or a linear combination across different value functions (i.e., $\overline{V}_{n2}$).

We must approximate $g(z, V)$ as a function of z and V. There are several ways to do this. One option is to use interaction terms between z and V and create a lower and upper envelope to $g(z, V)$ using planes. An alternative approach is to exploit the fact that V appears multiplicative to z, similar to the case of product prices in the Nash-Bertrand with discrete choice demand example.[34]

4.3.1 Numerical example for full dynamic learning with forgetting game.

Continuing the previous example, we set $\beta = 1/1.05$ and the forgetting parameter $\delta = (0.0, 0.0275, 0.1218)$. We bound the equilibrium policy and value functions for every state of the game, following the algorithm described in Appendix A.3. As we commented in the previous section and explain in the appendix, this algorithm iterates over the state space, solving a sequence of simplified versions of the program set up in Equation 5, which only includes a state-iteration dependent subset of the piece-wise linear relaxed conditions. Table 6 presents the bounds for $M = 8$.

As expected, bounds are very narrow for $\delta = 0$ since it has a unique equilibrium. For positive forgetting values, it is no longer possible to show that the equilibrium is unique. In such a case, the bounds can alert about the likelihood of multiplicity. We explore two cases. First, we take $\delta = 0.0275$. In this case, the bounds are also narrow. Thus, although we cannot formally prove equilibrium uniqueness under $\delta = 0.0275$, our bounds act as numerical proof.[35] On the other hand, the bounds on prices and values for $\delta = 0.1218$ are substantially wider.[36] It is indeed a forgetting value for which multiple equilibria arise. In Appendix B.3.1, we show that the game with eight states has at least two equilibria when $\delta = 0.1218$, directly solving the

[33] See appendix in Besanko et al. (2010) for a derivation. Note that we use a slightly different notation.

[34] Recall that Gurobi currently deals with quadratic interactions internally, solving quadratic non-convex problems globally.

[35] In any case, even if there are multiple equilibria for this forgetting parameter, our bounds indicate its implications are inconsequential.

[36] Readers familiar with Besanko et al. (2010) work might have noted that they find multiple equilibria for $\delta = 0.0275$ and do not find any multiplicity for $\delta = 0.12178$. Observe, however, that the dimension of our state space is lower than theirs. Moreover, as they noted, multiplicity depends on the relative sizes of forgetting and the state space, not their absolute values: *"We caution the reader that the absence of persistent asymmetries for low forgetting rates δ in Figure 3 may be an artifact of the finite size of the state space (M= 30 in our baseline parametrization). Given $\delta = 0.01$, say, $\Delta(30) = 0.26$ and organizational forgetting is so weak that the industry is sure to remain in or near state (30, 30). This minimizes bidirectional movements and restores the backward induction logic that underlies uniqueness of equilibrium for the that increasing M, while holding fixed δ, facilitates persistent asymmetries as extreme case of $\delta = 0$ (see s Proposition 3)." (Besanko et al., 2010, pp. 480).*

non-linear system of equations.[37]

Table 6: Counterfactual bounds on prices for the dynamic case

	$\delta = 0.0$	$\delta = 0.0275$	$\delta = 0.12178$
Average bounds $[\underline{p}_1, \bar{p}_1]$	[8.90, 8.93]	[8.84, 8.87]	[4.90, 9.93]
Iterated Prices p_1	8.92	8.87	7.09
Average $(\bar{p}_1 - \underline{p}_1)$	0.03	0.03	5.03
Maximum $(\bar{p}_1 - \underline{p}_1)$	0.06	0.08	14.14
Average bounds $[\underline{V}_1, \bar{V}_1]$	[19.88, 19.93]	[19.48, 19.55]	[6.70, 22.40]
Iterated Values V_1	19.93	19.55	11.31
Average $(\bar{V}_1 - \underline{V}_1)$	0.05	0.07	15.70
Maximum $(\bar{V}_1 - \underline{V}_1)$	0.11	0.14	23.40

Notes: Parameters are as in Table 4, with $\beta = 1/1.05$. The table presents averages across M^2 states. The solution to the game using a value-function iteration approach is also reported.

Additionally, Table 6 compares bounds to the equilibrium found using the classic value-function iteration approach. The bounds successfully cover the equilibrium to which the iterative procedure converged, even when these bounds are very narrow, i.e., when the equilibrium is unique.[38] This result suggests that the information loss from the piece-wise approximation will likely be small.

Even though loose bounds, as those reported in Table 6 for $\delta = 0.1218$, might appear relatively uninformative, keep in mind that Table 6 only reports average bounds. However, analyzing individual bounds, i.e., state by state, can provide valuable information. For instance, Figure 6 shows price and value bounds for every state at $\delta = 0.1218$. This figure indicates that the type of equilibria encountered using homotopy methods can also be inferred from our bounds. Indeed, the figure suggests the existence of equilibria with very aggressive pricing behavior: low prices when knowledge is similar among firms and low prices also to defend the advantage off-diagonal, and "flat-pricing" equilibria that resemble the equilibrium without forgetting[39]. We say "suggest" because the bounds do not represent any particular equilibrium but the minimum and maximum values and prices that can be achieved in any equilibrium.

[37] Figure B.5 presents these equilibria's policy and value function

[38] This is true for every state, not just for the average, as reported in the table.

[39] See Figure B.7 for similar plots with $\delta = 0$; policies and values for $\delta = 0.0275$ are remarkably similar to the case with $\delta = 0$, and are omitted

Figure 6: Bounds on equilibrium prices and values state by state

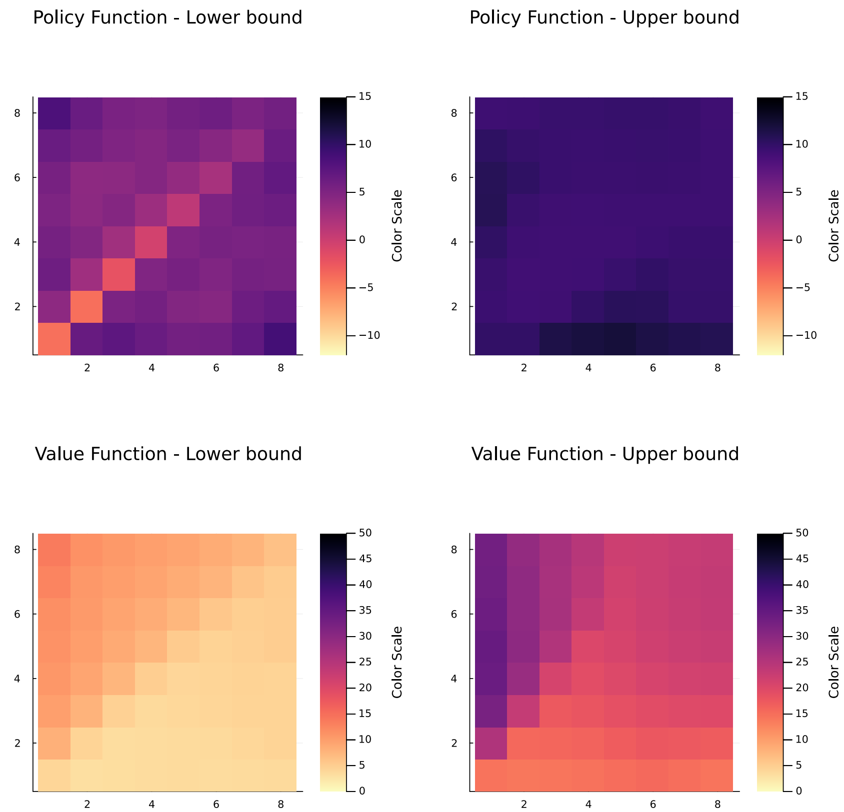

Notes: Parameters are as in Table 4, with $\beta = 1/1.05$ and $\delta = 0.12178$.

Interestingly, alternative solution methods miss out on the more standard flat pricing. The iterative solution only finds a locally stable equilibrium where prices feature a central trench.[40] Moreover, non-linear solvers only find two types of equilibria: equilibria with low prices throughout the diagonal (trenchy equilibrium) and flat-pricing with wells (low prices at $e = (0, 0)$), despite starting from multiple initial values (see Figure B.5). Hence, it is possible to characterize industry dynamics from equilibrium bounds and avoid misleading characterizations from other solution methods.

4.4 Equilibrium Selection

However, when multiplicity is large and bounds too wide, they can be indeed quite uninformative. Often, the researcher would like to put additional constraints on the game so that predictions on prices and/or firm behavior are less volatile. One advantage of the proposed method is that it is straightforward to incorporate additional equilibrium restrictions. As an advantage, such restrictions can be very explicit and motivated by features of the game or the institutional environment.

[40] Besanko et al. (2010) show Pakes and McGuire (1994)'s algorithm cannot find locally unstable equilibrium, which exists in the multiple equilibria case.

Consider the above dynamic game, where firms learn and reduce their costs. In regions where multiple equilibria arise, some firms compete very aggressively when they are close to each other, to the extent that they might set negative prices. There are also situations in which a firm, by reducing its costs, could trigger harsher competition. In those equilibria, there may be non-monotonicities. In particular, a firm might be better off (in NPV) in a state in which, *ceteris paribus*, its marginal costs are lower. It would prefer to increase its marginal costs ex-post if credible to the other firm.[41]

This example lends itself to illustrating how equilibrium selection works within our framework. That is, we can restrict equilibria to situations in which, *ceteris paribus*, firms are better off when their marginal costs are lower.[42] Using the proposed methodology, this can be easily embedded into the program by adding the following constraints:

$$\begin{aligned} V_1(e_1, e_2) &\geq V_1(e_1', e_2) \;\text{ if } C(e_1) < C(e_1'), \\ V_2(e_1, e_2) &\geq V_2(e_1, e_2') \;\text{ if } C(e_2) < C(e_2'). \end{aligned}$$

Failure to solve for the equilibrium conditions of the game under these additional restrictions implies that no equilibrium satisfies such a rule. On the other hand, if all equilibria satisfy this condition, adding extra information can still help "sharpen" the original bounds.[43]

Another surprising feature of learning-by-doing with forgetting is that firms may price below marginal cost to defend their position, even when they have reached the bottom of their learning curve. Naturally, they can also price below marginal cost while descending the learning curve, but this is not as surprising as there is an evident reward from doing so. Indeed, it is easy to see that there are no equilibria such that prices are always above marginal costs since, at the top of the learning curve, the price upper bound is actually below marginal cost.[44] We can also rule out equilibria with below-cost pricing at the bottom of the learning

[41] In Besanko et al. (2010) characterization, the trenchy equilibrium has this feature. Still, neither the flat nor the extra trenchy equilibria have it.

[42] A more strict restriction could be that the value function is monotonically increasing in the state, not just when marginal costs are lower.

[43] A similar—although not exactly equivalent—way to restrict the equilibrium is by imposing that dynamic prices are actually more competitive than static Nash-Bertrand for an equivalent demand. Following Besanko et al. (2010)'s characterization of prices:

$$p^*(\mathbf{e}) = c^*(\mathbf{e}) + \frac{\sigma}{1 - D_1^*(\mathbf{e})}$$

where $c^*(\mathbf{e}) = c(e_1) - \beta\phi^*(\mathbf{e})$ are *virtual* marginal costs, where $\beta\phi^*(\mathbf{e})$ represents the discounted prize for winning today's sale, that is

$$\phi^*(\mathbf{e}) = \bar{V}_1^*(\mathbf{e}) - \bar{V}_2^*(\mathbf{e})$$

Thus, we can restrict the equilibrium to be more competitive than static Nash-Bertrand by setting:

$$\bar{V}_1 - \bar{V}_2 \geq 0$$

[44] From a methodological point of view, it shows how we can use our approach to rule out equilibria.

curve, imposing additional constraints: [45][46]

$$P_1(e_1, e_2) \geq C_1(e_1), \quad \forall \quad e_1 \geq m$$
$$P_2(e_1, e_2) \geq C_2(e_2), \quad \forall \quad e_2 \geq m$$

There are many other alternative ways to select equilibria. For example, one could take a more Knightian approach and pick the equilibrium that minimizes the net present value at $t = 0$ when firms have the highest marginal cost. To do so, one could restrict the value function at $\{0, 0\}$ to be close to the lower bound found for that state using an unrestricted search, i.e., $V_1(0, 0) < \underline{V}(0, 0) + \epsilon$, refining the equilibrium iteratively if needed. Importantly, ϵ should be set large enough to allow for the potential conservativeness of the bounds.

The following section presents the results of restricting equilibria following. These bounds have been obtained following the algorithm described in Appendix A.3, and adding, in each case, the additional restrictions described above to a simplified version of optimization program 5.

4.4.1 Numerical example for a restricted dynamic game under multiplicity.

Table 7 presents the bounds of the game at $\delta = 0.1218$ under the two previously mentioned equilibrium selection restrictions. Column one presents the same bounds as in Table 6, corresponding to the unrestricted case. Column two restricts equilibria to be monotonic in its own marginal cost. Then, the last column presents the results from restricting prices to be above marginal costs at the bottom of the learning curve.

Table 7: Counterfactual bounds on prices for the dynamic case

	Unrestricted	Monotonicity	$p(e) > c(e) e \geq m$
Average bounds $[\underline{p}_1, \bar{p}_1]$	[4.90, 9.93]	[4.91, 9.93]	[6.75, 9.51]
Average $(\bar{p}_1 - \underline{p}_1)$	5.03	5.02	2.76
Maximum $(\bar{p}_1 - \underline{p}_1)$	14.14	14.14	9.37
Average bounds $[\underline{V}_1, \bar{V}_1]$	[6.70, 22.40]	[6.70, 22.37]	[7.66, 18.18]
Average $(\bar{V}_1 - \underline{V}_1)$	15.70	15.67	10.52
Maximum $(\bar{V}_1 - \underline{V}_1)$	23.40	23.37	15.33

[45] Imposing that at least one price is below marginal cost is somewhat more expensive from an optimization point of view, but can also be achieved with a combination of integer constraints, e.g.,

$$P_1(e_1, e_2) \leq C_1(e_1) + \lambda - \lambda U_1(e_1, e_2),$$
$$P_2(e_1, e_2) \leq C_2(e_2) + \lambda - \lambda U_2(e_1, e_2),$$
$$1 \leq \sum_{s,k} U_1(s, k) + U_2(s, k),$$

where λ is a large number, and $U_1, U_2 \in \{0, 1\}$.

[46] Note, these conditions do not increase the computational burden at all.

The monotonicity restriction does not seem to eliminate any equilibrium. This indicates, in principle, that all equilibria of the game satisfy the monotonicity assumption. Nevertheless, remember that our iterative procedure does not include all equilibrium equations at every step. Hence, some equilibria that do not meet the criteria may satisfy our elimination process. This can be easily solved by adding equilibrium constraints to each iteration step.

On the other hand, constraining prices to be above marginal costs at the bottom of the learning curve substantially narrows equilibrium bounds. In this case, both the lower bound on prices and the upper bound in values shrink. The first effect is by construction, as many prices within the lower bound violate our new constraints. The latter effect indicates that by raising credible threats to initiate a price war to defend its position, the leader increases the average value of the game.

Figure 7 shows again that the main implication of the restriction is to stop price wars to realize in equilibrium, lowering the leader's value. It eliminates the trenchy equilibrium, leaving only the flat pricing and flat pricing with wells as possible equilibrium strategies.

Figure 7: Bounds on equilibrium prices and values state by state

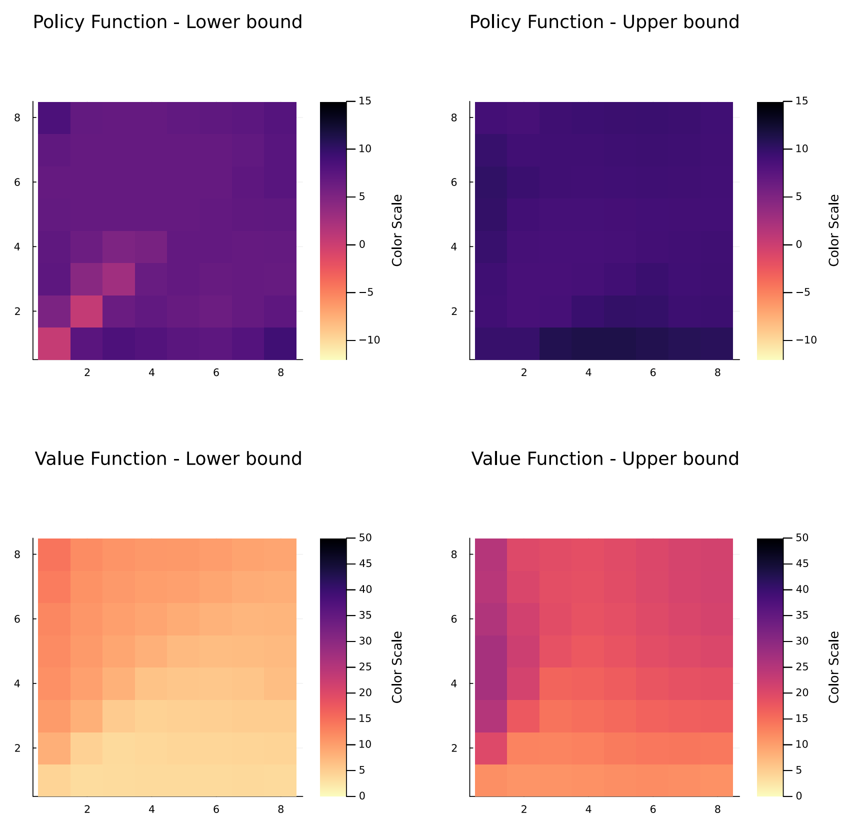

Notes: Parameters are as in Table 4, with $\beta = 1/1.05$ and $\delta = 0.12178$.

These examples illustrate how we can incorporate restrictions into our methodology to restrict equilibria plausibly and transparently.

5 Multi-Unit Auctions

The methods presented in this paper can be useful in the context of partially characterized equilibria, e.g. in situations in which only a set of necessary conditions for equilibrium can be characterized (but not necessary *and* sufficient conditions). There are several instances in which this might happen. One example is the case of multi-unit auctions, which are a form of supply function equilibrium and for which computation of counterfactuals is limited under rich cost structures and/or flexible distributions of uncertainty.[47] The inability to fully characterize supply function equilibria robustly has prevented the advancement of rich counterfactual experiments, which are often limited to VCG auctions in which bidders are truth-telling.

Multi-unit auctions are a form of supply function equilibria, where firms choose a schedule of prices and quantities, potentially a continuous object. However, as explained in Section 2, the proposed methodology is limited to situations in which the strategy space is finite. Therefore, this example also shows how the method can be applied when the strategy space is continuous by structuring the necessary conditions of the game in a way that is amenable to the methodology.

5.1 Model outline

Consider a multi-unit auction game in which players can choose a schedule of prices and quantities. A typical example are treasury bill auctions (Hortaçsu and McAdams, 2010; Kastl, 2011), or electricity markets (Wolak, 2003; Hortaçsu and Puller, 2008). At a given auction, an uncertain amount of units are being auctioned, A_s. This amount is uncertain before the auction. Hence, it is indexed with state $s = \{1, \ldots, S\}$.[48] Without loss of generality, we sort demand along states so that $A_s > A_{s-1}$.[49] Importantly, the dimensionality of demand states is assumed to be finite.

To implement the methodology, we focus on solving for the set of quantities q_{is} that each player $i = \{1, \ldots, I\}$ clears in equilibrium at each state s. One can interpret the resulting pairs of prices and quantities as representing the quantity-price schedule at endogenous breakpoints p_s, leaving the other points unspecified. Additional restrictions limit firms' strategy outside of the points on equilibrium prices. For example, as part of the rules in most multi-unit auctions, offers need to be monotonic. Therefore, as part of the equilibrium, it is required that,

$$p_s \leq p_{s+1}, q_{is} \leq q_{i,s+1}, \quad \forall s \in \{1, \ldots, S-1\}, \forall i.$$

There are often also price caps in multi-unit auctions, i.e., $p_s < \overline{p}, \forall s$. Finally, markets clear in equilibrium,

[47] Klemperer and Meyer (1989) discuss equilibrium computation under differentiability assumptions and unbounded support of the common uncertainty distribution. Under certain assumptions (e.g., linear strategies), Vives (2011) characterizes the equilibrium of the game even with richer information structures.

[48] In the Treasury market, this could be seen as residual uncertainty regarding the number of bidders with non-competitive tender bids, which shifts the inelastic demand for Treasury bills. In the case of electricity markets, there could be uncertainty about the amount of renewable power available at a given point in time.

[49] The demand curve can also have some slope, b, such that cleared demand equals $A_s - bp_s$. Without computational cost, one can also allow for rotations on the demand curve, as long as they do not make the ordering of states endogenous to p, i.e., as long as $b_s \leq b_{s-1}$.

Figure 8: Example of Candidate Supply Functions

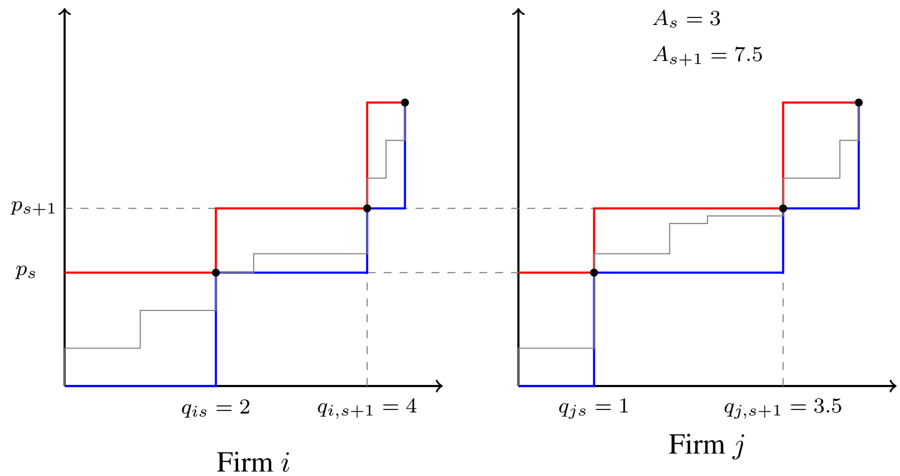

The pairs of offers $\{q_{is}, p_s\}$ and $\{q_{js}, p_s\}$ clear the market with demand $A_s = 3$ and price p_s. The pairs of offers $\{q_{i,s+1}, p_{s+1}\}$ and $\{q_{j,s+1}, p_{s+1}\}$ are clear the market when $A_{s+1} = 7.5$. The red envelope represents the upper bound to the supply function along these two points, whereas the blue envelope represents a lower bound. The gray line represents a candidate supply curve consistent with these envelopes.

i.e.,

$$\sum_i q_{is} = A_s, \quad \forall s.$$

Figure 8 depicts a candidate equilibrium for a set of supply functions that satisfy these bidding restrictions. Each of the two firms has a set of quantities and price pairs. At the price point p_s, firm i is offering $q_{is} = 2$ and firm j is offering $q_{js} = 1$, with $q_{is} + q_{js} = A_s$. Similarly, $q_{i,s+1} + q_{j,s+1} = A_{s+1}$. Because we do not solve for the full schedule of offers, we consider the upper (red) and lower (blue) envelopes on the offer curves, given monotonicity constraints. These offer points represent a potential equilibrium candidate for these sets of demands. They are valid because they satisfy the monotonicity constraint on offers, i.e., quantity schedules are increasing in price, and they clear the market. How can additional constraints be used to narrow down the equilibrium?

5.2 Necessary Conditions for Equilibrium

At any equilibrium schedule $\{p_s, \mathbf{q}_s\}_{s=1}^{S}$, no firm should want to deviate, taking the offer curve from the other firms as given. These deviations are similar to those considered in an estimation context by Hortaçsu and McAdams (2010). However, here, one is solving for the equilibrium strategies taking the costs as given, instead of inferring the costs from the equilibrium strategies.

Because the schedule as defined is discrete, we consider deviations to strategies taking a conservative version of other firms' schedule, i.e., taking the red and blue envelope as best and worst case scenarios to the supply function of a firm outside the points in the schedule that are solved for. Consider a situation where firm i offers q_{is} at state s with price p_s. For it to be weakly optimal, it needs to be the case that the firm does not want to lower its offered price and increase its quantity. Because the schedule of the other firms is not fully characterized, we consider a lower bound on deviation profits considering that, if the firm wants to produce more, it needs to offer at most p_{s-1}.[50] In such case, a lower bound deviation is given by,

$$p_s q_{is} - C(q_{is}) \geq p_{s-1}\Big(q_{i,s-1} + A_s - A_{s-1}\Big) - C\Big(q_{i,s-1} + A_s - A_{s-1}\Big),$$

where $C(.)$ is a weakly increasing cost function. The condition states that the firm should not be willing to increase its output at a lower price p_{s-1}. Note that given the auction rules and market clearing, $q_{is} \leq q_{i,s-1} + A_s - A_{s-1}$. Similarly, a firm should not be willing to increase its price at a given schedule and reduce its quantity. Assuming that a firm has to offer at least p_{s+1} to increase its quantity (as an upper bound), then,

$$p_s q_{is} - C(q_{is}) \geq p_{s+1}\Big(q_{i,s+1} + A_s - A_{s+1}\Big) - C\Big(q_{i,s+1} + A_s - A_{s+1}\Big),$$

where the firm reduces its output at a higher price, given the auction rules and market clearing imply $q_{is} \geq q_{i,s+1} + A_s - A_{s+1}$.

Because these are discrete deviations, they put less structure on the nature of the game. This has the advantage of not forcing the supply function schedules to satisfy certain properties (e.g., differentiability, linearity). One can also consider flexible cost functions, price caps, etc. As a disadvantage, these necessary conditions are not tight.

To implement the methodology, one can convert the above equilibrium conditions into linear constraints. Both revenues and costs are quadratic functions of prices and/or quantities. Following the methods from the previous examples, one can consider discretizing the price range and narrow down the envelopes at each state s iteratively to approximate revenues. For the cost function, one can approximate this single-dimensional function that depends on quantity with piece-wise linear envelopes. The cost function is already linear in the particular case of constant marginal costs. These constraints could be expanded, for example, to consider deviations of several steps at the same time.

5.3 Bounding Auction Costs

Given the above constraints, one can solve for the minimum and maximum procurement costs that can arise in a multi-unit auction, for example, in the context of electricity markets.[51] To do so, one can solve the

[50] Technically, to fully ensure that it wins all units at the margin, it needs to offer $p_{s-1} - \epsilon$.

[51] Equivalently, in the context of Treasury bills, one can find upper and lower bounds to auction revenues, an object that is often considered as the primal counterfactual of interest (Kang and Puller, 2008; Kastl, 2011).

following program,

$$
\begin{aligned}
\max_{\mathbf{p},\mathbf{q}} \text{ and } \min_{\mathbf{p},\mathbf{q}} \quad & \sum_s p_s A_s, \\
\text{s.t.} \quad & p_s > p_{s-1}, q_{is} > q_{i,s-1}, p_s \leq \overline{p}, \\
& \textstyle\sum_i q_{is} = A_s, \\
& p_s q_{is} - C(q_{is}) \geq p_{s-1}(q_{i,s-1} + A_s - A_{s-1}) - C(q_{i,s-1} + A_s - A_{s-1}), \\
& p_s q_{is} - C(q_{is}) \geq p_{s+1}(q_{i,s+1} + A_s - A_{s+1}) - C(q_{i,s+1} + A_s - A_{s+1}),
\end{aligned}
$$

where the non-linear constraints on revenues and marginal deviations are approximated using piece-wise linear constraints.

We consider an example related to electricity auctions. Two firms are competing in a multi-unit auction. Each has four power plants with capacity constraints and unit-specific marginal costs. Demand is uncertain and given by A. There is also a price cap, which constrains the maximum price that firms can set. We solve for the optimal strategies of the firms given a demand distribution $\{A_1, \ldots, A_S\}$. Uncertainty in demand comes from residual uncertainty about other firms and/or renewable production rather than demand, which is typically forecasted accurately. We assume that the two firms have substantial market power. Therefore, they might become pivotal, i.e., for some state s their production is needed to cover demand. We also assume that the lowest demand realization is above zero, i.e. $A_s > 0$, which makes sense in the context of electricity markets, as uncertainty is usually bounded.

A typical concern to policymakers in electricity markets is the expected cost of electricity to consumers, i.e., $\sum_s \omega_s p_s A_s$. Therefore, we look for bounds on the procurement costs of electricity, allowing the equilibrium concept to be quite flexible (i.e., only imposing the equilibrium constraints described above). Figure 9 shows the aggregate supply curves that generate the highest (red) and lowest (blue) procurement costs. One can see that even the scenario with the lowest procurement costs features a supply curve above marginal costs. The scenario with the highest procurement costs is below the price cap, except for the highest demand realization.

Figure 9: Bounds on Procurement Outcomes

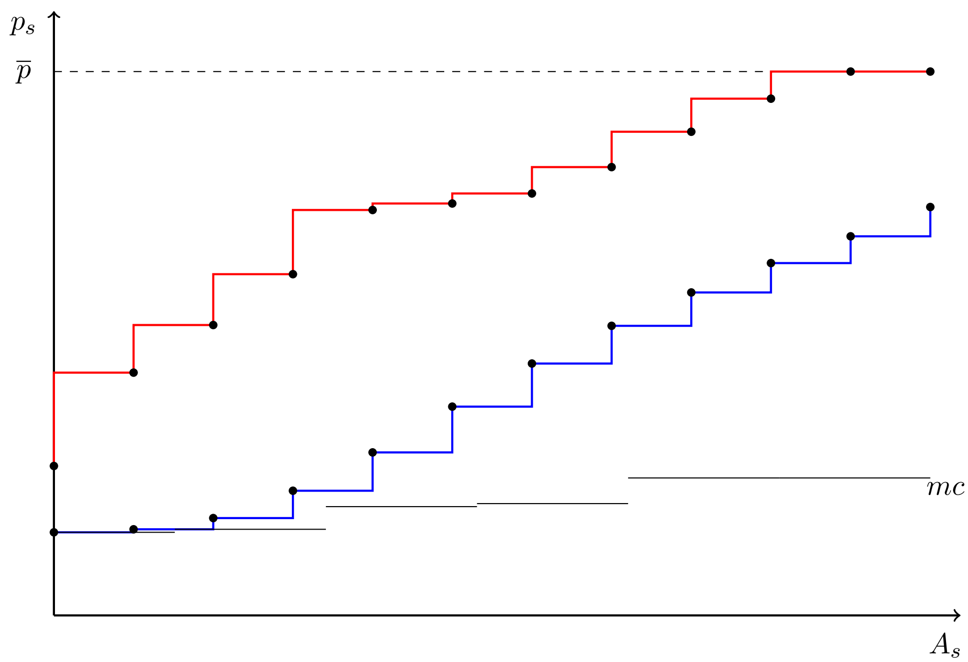

The red bound shows the supply curve that generates the highest procurement costs, whereas the blue curve shows the supply curve that generates the lowest procurement costs. These supply curves are consistent with necessary equilibrium conditions and provide valid conservative bounds to procurement costs. In this particular example, it is clear from the graph that procurement costs are bounded above marginal costs but also bounded below the price cap.

6 Conclusions

We have presented a new methodology to bound counterfactual outcomes. The basic idea is to minimize and maximize the counterfactual outcome of interest, subject to the outcome being consistent with a set of equilibrium constraints. The method relies on three guiding principles: non-iteration, linear optimization, and relaxation. These principles are intended to make the counterfactual bounds valid and robust, at the expense of being more conservative.

From a technical point of view, the method is implemented using optimization with integer variables and piece-wise linear constraints. This represents two innovations from the previous literature. First, we propose to use a relaxed constrained-optimization approach for equilibrium computation. Whereas constrained approaches have been typically used for estimation or computation of single-agent problems, we show how to use them for computation purposes in strategic settings. In particular, we show how it can help address the issue of multiple equilibria. Because the equilibrium conditions are expressed as a set of constraints, the objective function can be used to bound the counterfactual outcome of interest. Second, the method emphasizes using robust equilibrium *bounds* instead of exact equilibrium equations. By relying on bounding techniques and mixed-integer formulations, one can guarantee the validity of the bounds.

To show the appropriateness of the method, we have presented three examples in which counterfactual simulations might be difficult or multiplicity of equilibria might arise. First, we presented games of static price competition with multiple equilibria and showed how counterfactual outcomes can be computed. Second, we explored the implementation of the methodology in Markov dynamic games, proposing an alternative approach to compute equilibria that complements the literature on homotopy methods. Finally, we showed how to use the method to compute flexible equilibrium bounds when firms compete in supply functions, an area in which we were still limited in the counterfactuals that we could perform.

References

Abbring, J., Campbell, J., Tilly, J., and Yang, N. (2016). Very Simple Markov-Perfect Industry Dynamics. Discussion Paper 11069, CEPR.

Aguirregabiria, V. (2012). A method for implementing counterfactual experiments in models with multiple equilibria. *Economics Letters*, 114(2):190–194.

Aguirregabiria, V. and Mira, P. (2012). Structural Estimation and Counterfactual Experiments in Games when Data come from Multiple Equilibria. Technical report, mimeo.

Aksoy-Pierson, M., Allon, G., and Federgruen, A. (2013). Price competition under mixed multinomial logit demand functions. *Management Science*, 59(8):1817–1835.

Bajari, P., Benkard, C. L., and Levin, J. (2007). Estimating dynamic models of imperfect competition. *Econometrica*, 75:1331–70.

Bajari, P., Hong, H., and Ryan, S. P. (2010). Identification and estimation of a discrete game of complete information. *Econometrica*, 78(5):1529–1568.

Benkard, C. L. (2004). A Dynamic Analysis of the Market for Widebodied Commercial Aircraft. *Review of Economic Studies*, 71:581–611.

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile Prices in Market Equilibrium. *Econometrica*, 63(4):841–90.

Berry, S., Levinsohn, J., and Pakes, A. (1999). Voluntary Export Restraints on Automobiles: Evaluating a Trade Policy. *American Economic Review*, 89(3):400–430.

Besanko, D., Doraszelski, U., Kryukov, Y., and Satterthwaite, M. (2010). Learning-by-Doing, Organizational Forgetting, and Industry Dynamics. *Econometrica*, 78(2):453–508.

Bresnahan, T. F. and Reiss, P. C. (1990). Entry in Monopoly Markets. *The Review of Economic Studies*, 57(4):pp. 531–553.

Caplin, A. and Nalebuff, B. (1991). Aggregation and Imperfect Competition: On the Existence of Equilibrium. *Econometrica*, 59(1):25–59.

de Boor, C. (2001). *A Practical Guide to Splines. Revised Edition.* Springer-Verlag New York Inc.

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

Dubé, J.-P., Fox, J. T., and in Su, C. (2012). Improving the Numerical Performance of Static and Dynamic Aggregate Discrete Choice Random Coefficients Demand Estimation. *Econometrica*, 80(5):2231–2267.

Dube, J.-P., Hitsch, G. J., and Rossi, P. E. (2010). State dependence and alternative explanations for consumer inertia. *The RAND Journal of Economics*, 41(3):417–445.

Echenique, F. and Komunjer, I. (2007). Testing Models With Multiple Equilibria by Quantile Methods. Technical Report Working Paper 1244R, Division of Humanities and Social Sciencies, Caltech.

Ericson, R. and Pakes, A. (1995). Markov-Perfect Industry Dynamics: A Framework for Empirical Work. *The Review of Economic Studies*, 62(1):pp. 53–82.

Farias, V., Saure, D., and Weintraub, G. (2012). An approximate dynamic programming approach to solving dynamic oligopoly models. *RAND Journal of Economics*, 43:253–282.

Fershtman, C. and Pakes, A. (2012). Dynamic games with asymmetric information: A framework for empirical work. *The Quarterly Journal of Economics*, 127(4):1611–1661.

Fowlie, M. L., Reguant, M., and Ryan, S. P. (2014). Market-Based Emissions Regulation and Industry Dynamics. *Journal of Political Economy, forthcoming*.

Grennan, M. and Town, R. (2015). Regulating innovation with uncertain quality: Information, risk, and access in medical devices. Working Paper 20981, National Bureau of Economic Research.

Guerre, E., Perrigne, I., and Vuong, Q. (2000). Optimal Nonparametric Estimation of First-Price Auctions. *Econometrica*, 68(3):525–574.

Haile, P. A. and Tamer, E. (2003). Inference with an Incomplete Model of English Auctions. *Journal of Political Economy*, 111(1):1–51.

Hanson, W. and Martin, K. (1996). Optimizing multinomial logit profit functions. *Management Science*, 42(7):992–1003.

Hortaçsu, A. and McAdams, D. (2010). Mechanism Choice and Strategic Bidding in Divisible Good Auctions: An Empirical Analysis of the Turkish Treasury Auction Market. *Journal of Political Economy*.

Hortaçsu, A. and Puller, S. L. (2008). Understanding Strategic Bidding in Multi-unit Auctions: A Case Study of the Texas Electricity Spot Market. *RAND Journal of Economics*, 39(1):86–114.

Jia, P. (2008). What Happens When Wal-Mart Comes to Town: An Empirical Analysis of the Discount Retailing Industry. *Econometrica*, 76(6):1263–1316.

Jofre-Bonet, M. and Pesendorfer, M. (2003). Estimation of a Dynamic Auction Game. *Econometrica*, 71(5):1443–1489.

Judd, K. L., Renner, P., and Schmedders, K. (2012). Finding All Pure-Strategy Equilibria in Games with Continuous Strategies. *Quantitative Economics*, 3(2):289–331.

Kang, B.-S. and Puller, S. L. (2008). The Effect of Auction Format on Efficiency and Revenue in Divisible Good Auctions: A Test Using Korean Treasury Auctions. *The Journal of Industrial Economics*, 56(2):290–332.

Kastl, J. (2011). Discrete Bids and Empirical Inference in Divisible Good Auctions. *Review of Economic Studies*, 78(3):974–1014.

Kellogg, R. (1976). Uniqueness in the schauder fixed point theorem. *Proceedings of the American Mathematical Society*, 60(1):207–210.

Klemperer, P. D. and Meyer, M. A. (1989). Supply Function Equilibria in Oligopoly under Uncertainty. *Econometrica*, 57(6):1243–77.

Konovalov, A. and Sándor, Z. (2010). On price equilibrium with multi-product firms. *Economic Theory*, 44(2):271–292.

Leslie, P. (2004). Price Discrimination in Broadway Theater. *RAND Journal of Economics*, 35(3):520–541.

Maskin, E. and Tirole, J. (1988). A Theory of Dynamic Oligopoly, I: Overview and Quantity Competition with Large Fixed Costs. *Econometrica*, 56:549–69.

Milgrom, P. and Roberts, J. (1990). Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities. *Econometrica*, 58(6):1255–77.

Mizuno, T. (2003). On the existence of a unique price equilibrium for models of product differentiation. *International Journal of Industrial Organization*, 21(6):761–793.

Nevo, A. (2000a). A Practitioner's Guide to Estimation of Random-Coefficients Logit Models of Demand. *Journal of Economics & Management Strategy*, 9(4):513–548.

Nevo, A. (2000b). Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry. *RAND Journal of Economics*, 31(3):395–421.

Nocke, V. and Schutz, N. (2018). Multiproduct-firm oligopoly: An aggregative games approach. *Econometrica*, 86(2):523–557.

Pakes, A. and McGuire, P. (1994). Computing markov-perfect nash equilibria: Numerical implications of a dynamic differentiated product model. *The Rand Journal of Economics*, pages 555–589.

Pakes, A. and McGuire, P. (2001). Stochastic algorithms, symmetric markov perfect equilibrium, and the 'curse'of dimensionality. *Econometrica*, 69(5):1261–1281.

Pakes, A., Ostrovsky, M., and Berry, S. (2007). Simple estimators for the parameters of discrete dynamic games (with entry/exit examples). *The RAND Journal of Economics*, 38(2):373–399.

Petrin, A. (2002). Quantifying the Benefits of New Products: The Case of the Minivan. *Journal of Political Economy*, 110(4):705–729.

Ryan, S. P. (2012). The cost of environmental regulation in a concentrated industry. *Econometrica*, 80(3):1019–61.

Su, C. and Judd, K. L. (2012). Constrained Optimization Approaches to Estimation of Structural Models. *Econometrica*, 80(5):2213–2230.

Sweeting, A. (2009). The strategic timing incentives of commercial radio stations: An empirical analysis using multiple equilibria. *RAND Journal of Economics*, 40(4):710–742.

Uetake, K. and Watanabe, Y. (2014). Entry by Merger: Estimates from a Two-Sided Matching Model with Externalities. Discussion Paper 2188581, SSRN.

Vives, X. (1999). *Oligopoly pricing: old ideas and new tools*. MIT press.

Vives, X. (2011). Strategic supply function competition with private information. *Econometrica*, 79(6):1919–1966.

Weintraub, G. Y., Benkard, C. L., and Roy, B. V. (2008). Markov Perfect Industry Dynamics with Many Firms. *Econometrica*, 76(6):1375–1411.

Wolak, F. A. (2003). *Identification and Estimation of Cost Functions Using Observed Bid Data: An Application to Competitive Electricity Markets*, chapter 4, pages 133–169. Cambridge University Press.

A Computational Appendix

A.1 MIP relaxation of equilibrium constraints

This section illustrates how to modify a non-linear optimization program, relaxing equilibrium constraints. We use the discrete type Nash-Bertrand pricing game as an example. The main objective is to move from a non-linear system of equations to a mixed-integer linear programming (MIP) problem, for which we can find global optima. We proceed in two steps. First, we construct a piece-wise upper and lower envelope of the non-linear function we wish to approximate. Second, we introduce auxiliary variables and restrictions that transform the non-linear equations into the appropriate linear equality, inequality, and integer restrictions. Finally, we suggest how to deal with quadratic interactions.

A.1.1 Auxiliary variables and restrictions.

We discuss the second step first, assuming we have computed piece-wise linear envelopes. We re-write the optimization program in 3.1 taking $\gamma_{ij} = ln(s_{ij})$ as a variable, and $f(x) = exp(x)$ as the function to be approximated, with $s_{ij} \in (\underline{f}(\gamma_{ij}), \bar{f}(\gamma_{ij}))$.

Suppose we have already built lower and upper bounds approximation of the exponential function. That is, we have two collections of linear coefficients $\{\beta^g_{lb}, \beta^g_{ub}\}_{g\in\{1,..,G+1\}}$, where $G+1$ is the number of nodes used in the approximation, the subscript refers to the upper or lower envelope and the superscript refers to the slope of the approximation between nodes g and $g-1$ for $g > 1$, while $\beta^1_{lb}, \beta^1_{ub}$ are the constants in each approximation.

If we are approximating the non-linear function $f(x)$ using G+1 nodes: $\{X[1], ..., X[G+1]\}$, then we define auxiliary variables $\{gZ^g_{ij}, uZ^g_{ij}\}_{i\in 1,..,I, j\in 1,..,J, g\in 1,..,G}$, where $gZ^g_{ij} \leq X[g+1] - X[g]$. Also define $\{uZ^g_{ij}\}_{i\in 1,..,I, j\in 1,..,J, g\in 1,..,G}$, which are binary variables. Finally, define another set of variable $\{s_{ij}\}_{i\in 1,..,I, j\in 1,..,J}$, which are bounded by our lower and upper envelope approximations of $f(x)$ evaluated at γ_{ij}. We specify $\{gZ_{ij}, uZ_{ij}\}_{g\in 1,...,G}$ and s_{ij} using the following constraints for every i and j

ALGORITHM 1

Auxiliary variable specification for the MIP transformation of a non-linear optimization program.

$\gamma_{ij} = X[1] + \sum_{g\in 1,...,G} gZ^g_{ij}$
for $g \in 1, ..., G$ **do**
  $gZ^g_{ij} \leq uZ^g_{ij} \times (X[g+1] - X[g])$
  if $g > 1$ **then**
    $gZ^g_{ij} \geq uZ^g_{ij} \times (X[g] - X[g-1])$
    $uZ^{g-1}_{ij} \leq uZ^g_{ij}$
  end if
end for
$s_{ij} \leq \beta^1_{up} + \sum_{g\in 1,..,G} \beta^g_{ub} gZ^g_{ij}$
$s_{ij} \geq \beta^1_{lb} + \sum_{g\in 1,..,G} \beta^g_{lb} gZ^g_{ij}$

Although all variables (main and auxiliary) will be solved simultaneously, intuitively, the role of $\{gZ_{ij}^g, uZ_{ij}^g\}_{g\in G}$ is to determine which pieces of the approximation are active and to which extent. They ensure we evaluate the piece-wise linear envelopes at the right point, that is, at γ_{ij}. Take, for instance, the case with $G+1=5$, and γ_{ij} in equilibrium lying within $\in (X[3], X[4])$, in equilibrium. Then, restrictions in Algorithm 1 suggest that auxiliary variables can be represented as in Figure A.1

Figure A.1: Evaluating piece-wise linear functions at desired points.

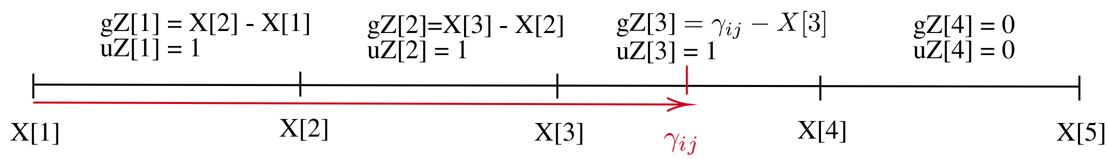

Then, we can easily evaluate the envelopes at the desired point by applying the corresponding linear coefficients to each interval $gZ[g]$. These linear coefficients, in turn, are determined by the parameterized piece-wise envelopes of the exponential function: $\{\beta_{lb}^g, \beta_{ub}^g\}_{g\in 1,..,G+1}$.

$$s_{ij}^{ub} = \beta_{up}^1 + \sum_{g\in 1,..,G} \beta_{ub}^g gZ_{ij}^g$$
$$s_{ij}^{lb} = \beta_{lb}^1 + \sum_{g\in 1,..,G} \beta_{lb}^g gZ_{ij}^g$$

Finally, the program enforces that s_{ij} is bounded by the lower and upper envelop approximations, i.e. $s_{ij} \in (s_{ij}^{lb}, s_{ij}^{up})$. This is represented in Figure A.2.

Figure A.2: Bounds on market shares, conditional on its logarithmic value.

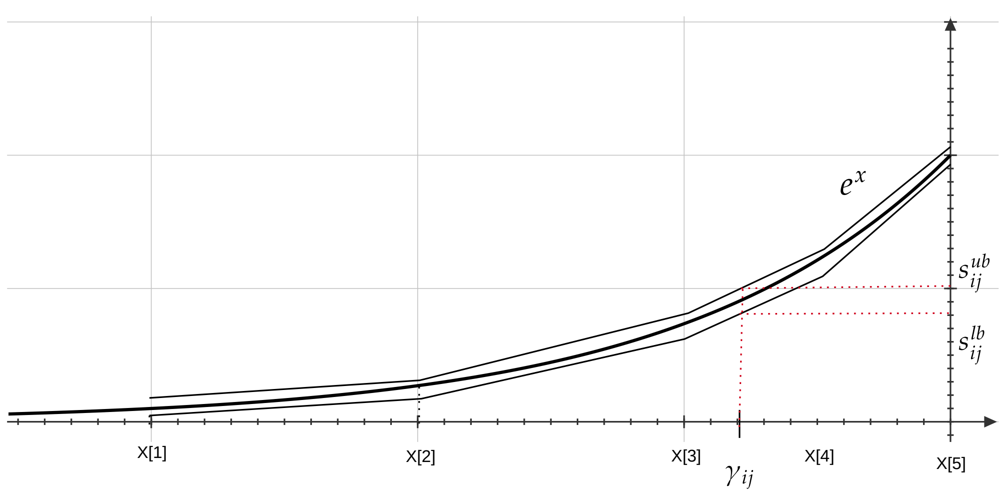

Notes: Example with G+1 =5. Nodes X[j] need not be equidistant. Indeed, node positioning can increase the efficiency of the method.

A.1.2 Piece-wise linear approximation.

In the previous section, we showed how to use piece-wise linear envelopes to transform a non-linear optimization program into a MIP problem. Now, we describe how to construct such envelopes for any function $f(x)$, using the exponential as an example.

First, we must bound the support of the approximated bounds. For instance, if we are trying to approximate shares through the exponential function, it is natural to assume that the support will be between a large negative number (shares close to 0) and 0 (shares close to 1). These bounds determine our linear envelopes' lowest and highest nodes, that is, $X[1]$ and $X[G+1]$, respectively. Next, we choose the number of nodes and how to position them over the support. As discussed in the main text, we can either follow some rule of thumb or use sophisticated techniques as optimal knot placement routines -see (de Boor, 2001) for a discussion in the context of splines. There is, of course, a tradeoff between the computational burden of including more knots and the precision of the approximations. To be clear, the computational cost does not come from approximating the function over more nodes. It arises because it eventually translates into adding more integer variables in the optimization routine. Proposition 4 suggests a way to circumvent this tradeoff partially.

Once we have chosen the number of nodes and their position in the support, we use constrained OLS to find the slope of each approximating piece. That is, the upper bound will be the piece-wise linear function that minimizes the distance to the function but is weakly larger at every point. At the same time, the lower bound also minimizes the distance but is weakly smaller at every point.

As an illustration, suppose we want to construct lower and upper envelopes for the exponential function over five nodes, between $X[1]$ and $X[5]$, as shown in Figure A.3. Define N points within the support $[X[1], X[5]]$, where N is a large integer. Moreover, evaluate the function to be approximated in each one of these points, that is, $Y_i = exp(X_i)$, where $X_i \in [X[1], X[5]] \quad i \in 1, ..., N$. Lets' $\hat{Y}_{lb}$ and $\hat{Y}_{ub}$ be the lower and upper envelope of $exp()$ at X_i. Then, in each case, we minimize the prediction error $\hat{Y}_i - Y_i$, subject to the error being positive in the case of the upper bound and negative for the lower bound.

Note, however, that to compute $\hat{Y}_{lb}$ and $\hat{Y}_{ub}$, we must define which pieces are active at each point X_i. Figure A.3 illustrates this procedure. For instance, for the point ZG, the interval $X[2] - X[1]$, $X[3] - X[2]$ and $X[4]-X[3]$ are active, wile $X[5]-X[4]$ is not. Thus, for each point X_i, we can define an auxiliary set of variables $d(X_i)$ of dimension G, which determines the influence of each linear coefficient in the prediction. In our example

$$d(ZG) = \begin{pmatrix} X[2] - X[1] \\ X[3] - X[2] \\ ZG - X[3] \\ 0 \end{pmatrix}$$

Next, we define the predicted piece-wise envelope value at X_i for the lower and upper envelopes as

$$\hat{Y}_{lb,ub}[X_i] = \beta^1_{lb,ub} + \sum_{g \in 2,...,G+1} d(X_i)[g]\beta^g_{lb,ub}$$

which is shown in Figure A.3 for the case of $X_i = ZG$

Figure A.3: Piecewise approximation of an exponential function. G+1 = 5

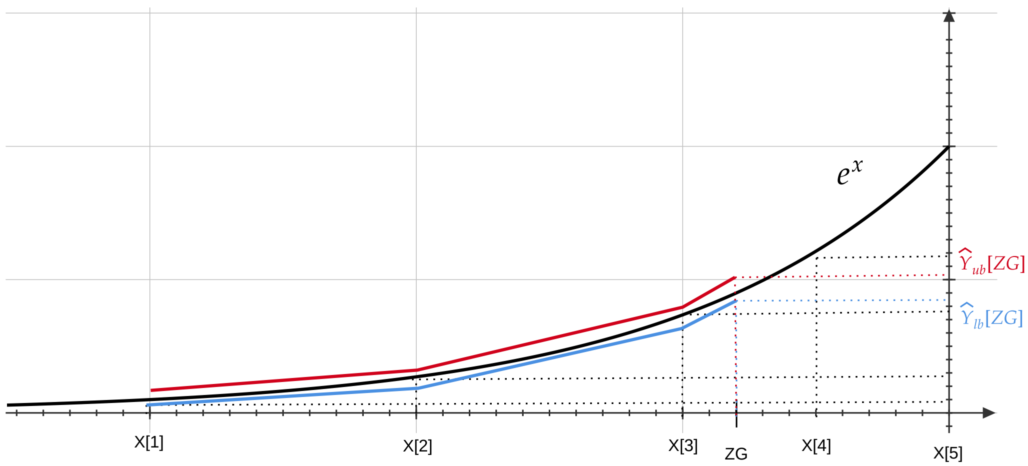

Notes: . Nodes X[j] need not be equidistant. Indeed, node positioning can increase the efficiency of the method.

Finally define regression errors $\epsilon_i = Y_i - \hat{Y}_i$ and compute $\{\beta^g_{lb}, \beta^g_{ub}\}_{g \in G+1}$ by solving

$$\beta_{lb} = \text{argmin}_\beta \sum_{i \in 1,..,N} \epsilon_i(\beta)$$
$$\text{s.t} \quad \epsilon_i \geq 0 \quad \forall i \in 1, ..., N$$

and,

$$\beta_{ub} = \text{argmax}_\beta \sum_{i \in 1,..,N} \epsilon_i(\beta)$$
$$\text{s.t} \quad \epsilon_i \leq 0 \quad \forall i \in 1, ..., N$$

A.1.3 Dealing with quadratic interactions.

In Section 3.1, we observed that approximating the shares and shares interactions is not enough to linearize the necessary conditions since markups interact with shares in firms' first-order conditions. Optimization software, such as Gurobi, can deal with non-convex quadratic interactions up to global optimality. Therefore, a simple way to circumvent the interaction problem is to write the quadratic program and let the optimizer solve it up to global optimality.

However, such an approach might imply giving up flexibility and can be computationally intensive. A straightforward approach to dealing with quadratic interaction is to partition one of the interacted variables into a grid of points and impose that the interaction must lie within the upper and lower bound of the points of the bins to which the actual variable belongs.

Suppose we partition products' markups into a grid of $K+1$ points $\{mG_{jk}\}_{k=1}^{K+1}$ and construct K bins such that the lower and upper bounds of markup j in bin k is $[mG_{kj}, mG_{k+1j}]$. Then, define auxiliary variables $um_{jk} \in \{1, ..., K\}$ such that

$$\sum_{k=1}^{K} um_{jk} = 1$$

which indicates that only one bin can be active, or those markups must lie within only one bin.

Moreover, the actual markup must be bounded by the lower and upper bounds of the active bin. Let m_j denote the markup variable in our optimization program, then

$$m_j \geq \sum_{k=1}^{K} um_{jk} \times mG_{jk}$$
$$m_j \leq \sum_{k=1}^{K} um_{jk} \times mG_{j,k+1}$$

Let γ_{ij} be the previously constructed auxiliary variable representing the interaction between market shares that is $\underline{\gamma}(s_{ij}(1-s_{ijt})) \leq \gamma_{ij} \leq \bar{\gamma}(s_{ij}(1-s_{ijt}))$ and M be a large positive constant. Then, we can bound the non-linear term in firms' FOC by imposing the following K conditions

$$\alpha_i s_{ij}(1-s_{ij})m_j \leq \alpha_i \gamma_{ij} \times mG_{j,k+1} + M - M \times um_{jt} \quad \forall k \in \{1, ..., K\}$$
$$\alpha_i s_{ij}(1-s_{ij})m_j \geq \alpha_i \gamma_{ij} \times mG_{j,k} - M + M \times um_{jt} \quad \forall k \in \{1, ..., K\}$$

A.2 Automatic MIP relaxation: implementation in Gurobi

Section A.1 explained how to use auxiliary variables and constraints to transform—relax—a non-linear equilibrium problem into a MIP optimization program. Although this procedure can be extrapolated to non-linear environments, the number of functions and auxiliary variables we must introduce can quickly become intractable, making code error-prone. This is the case, for instance, in Section 3.2's conditional heteroskedastic example. In this section, we describe how to use the popular solver to automatically transform the original non-linear optimization program into a MIP program, which can then be solved up to global optimality.*Gurobi*[52]

Gurobi allows users to define general, non-linear constraints, which are automatically transformed into piece-wise linear constraints. The advantage is that the user does not need to go into the *"esoteric details of how to model these relationships in terms of more fundamental constraints of MIP"*.[53] However, this is not exactly what our method does. We propose constructing piece-wise envelopes around the equilibrium constraints instead of piece-wise linear approximations of the non-linear constraints.

[52] This functionality is available for versions 9.0 and above.

[53] https://www.gurobi.com/documentation/9.5/refman/constraints.html#subsubsection:SOSConstraints

Nevertheless, we can still construct such envelopes leveraging the attributes of *Gurobi*'s *Function Constraints*. When defining a *Function Constraint*, the user can determine the number of pieces (*FuncPieces*), the length of the pieces (*FuncPieceLength*), or the approximation error (*FuncPieceError*). Furthermore, it allows the user to set whether the approximation is an underestimate of the function (*FuncPieceRatio* = 0.0) and overestimate (*FuncPieceRatio* = 1.0) or something in between. Therefore, by setting *FuncPieceRatio* = 0.0, we can automatically construct the lower piece-wise envelope of the function, and by setting *FuncPieceRatio* = 1.0, we build the upper envelope.

In our Nash-Bertrand game with conditional heteroskedastic logit demand, we need to construct piece-wise linear envelopes for the relation between log-shares and shares, similar to A.1. However, we must also bound the conditional heteroskedastic polynomial function and its derivative and relax multiple quadratic interactions. In all cases we can create auxiliary variables $\{(s_{ij}^{lb}, s_{ij}^{ub}), (g^{lb}, g^{ub}), (dg^{lb}, dg^{ub})\}$ defined by the Function Constraints under the appropriate values of *FuncPieceRatio*. Then, we define a set of linear constraints that determine that the individual shares $\{s_{ij}\}$, the variance g, and its derivative, dg, must be within these piece-wise linear envelopes.

$$
\begin{aligned}
s_{ij}^{lb} \leq s_{ij} \leq s_{ij}^{ub} \\
g^{lb} \leq g \leq g^{ub} \\
dg^{lb} \leq dg \leq dg^{ub}
\end{aligned}
$$

A.2.1 A comparison between Gurobi's automatic MIP transformation and manual implementations.

This section explores the comparative analysis of the welfare bounds for the multi-peak problem between Gurobi's automatic Mixed Integer Programming (MIP) transformation and manual implementations. The comparison primarily considers three dimensions: the internal handling of interactions in the non-convex problem, the number of nodes used in the approximation (a factor directly affecting the computational time), and the allowed error in the Gurobi implementation. We evaluate them at $\alpha = -2.524$, where equilibrium is unique, and consumer surplus equals 5.1248.

From Table 8, it can be observed that there is a significant variation in the execution times and the sizes of bounds across different models and the number of approximation nodes. This table helps us understand each model's performance, approximation precision, and the effects of varying the approximation nodes and tolerances.

The model without quadratic interactions is a purely MIP problem, solved in one shot with $\{5, 50, 100\}$ nodes. The model with cubic interactions in the second column allows Gurobi to handle all non-linear variable interactions internally while still finding a global solution. Third, the iterated approach constructs bounds iteratively. That is, it solves the problem several times and, at each iteration, uses the bounds built in the previous one. Finally, the full automatic implementations (1) and (2) are run using Gurobi general constraints to create piece-wise linear envelopes of the exponential function. In these cases, we restrict the FuncPieceError argument in Gurobi (maximum absolute approximation error) to 1e-3 and 1e-4, respectively.

Table 8: Execution times (in seconds) and bounds sizes for different models and sizes of approximation nodes.

		Model				
Nodes	Error PWL	No Quadratic Interactions	Cubic Interactions	Iterated Approximation	Automatic Bounds (1)	Automatic Bounds(2)
5		[0.96, 5.45]	[0.98, 5.44]	[5.11, 5.12]	*	*
		1.8	0.8	22.0	*	*
50		[2.15, 5.37]	[2.34, 5.27]	*	*	*
		32.0	19.2	*	*	*
100		[5.03, 5.25]	[5.12, 5.20]	*	*	*
		106.9	47.9	*	*	*
	1e-3	*	*	*	[5.10, 5.14]	*
		*	*	*	46.0	*
	1e-4	*	*	*	*	[5.12, 5.13]
		*	*	*	*	862.7

Automatic Bounds (1) and (2) restrict the FuncPieceError (the maximum absolute approximation error) to 1e-3 and 1e-4, respectively. The "Iterated Approximation" model allows for cubic interactions.

The results indicate that with five nodes, the "No Quadratic Interactions" and "Cubic Interactions" models produced similarly broad bound. Recall that, in this case; the equilibrium is unique. Thus, the size of the bounds is entirely due to computational limitations. Both models require approximately 100 nodes to achieve informative bounds. In this case, it should be noted that the "Cubic Interactions", which let Gurobi internally handle non-linear interactions, showed superior computational efficiency compared to the "No Quadratic Interactions" model.

The results with the specified errors (1e-3) demonstrate the benefits of utilizing automatic bounds in Gurobi. It provides a more precise bounding range than the previous models without increasing the run time. On the other hand, the automatic bound with an error equal to 1e-4 illustrates that considering Gurobi's automatic bounds transformation or manual implementations ultimately depends on the tradeoff between computational time and the desired precision of the solution since narrowing down the bounds can be very costly.

Finally, the iterated approach performs better than any other model in both dimensions. It obtains narrower bounds in a shorter period. Whether this approach dominates the other for different applications remains an open question. We have usually found it to perform at least comparably well to other approaches in various cases.

To sum up, the choice of the model ultimately depends on the specifics of the problem at hand and the resources available for computation. In any case, Gurobi's automatic MIP transformation provides a robust, precise, and off-the-shelf alternative when considering complex non-linear optimization problems.

A.3 An stochastic and iterative algorithm to compute equilibrium bounds

Suppose the objective is to find the minimum and maximum price and value that can be sustained in equilibrium at each state. We could solve the following program $M^2 \times 2 \times 2$ times.

$$\begin{aligned} \max_{\mathbf{V}^*,\mathbf{p}^*} \text{ and } \min_{\mathbf{V}^*,\mathbf{p}^*} \quad & V^*(e) \text{ or } P^*(e) \\ \text{s.t.} \quad & \underline{\mathbf{F}}(\mathbf{V}^*,\mathbf{p}^*) \le 0 \le \overline{\mathbf{F}}(\mathbf{V}^*,\mathbf{p}^*), \end{aligned} \tag{6}$$

This is, nevertheless, a complicated system to solve, which increases with the dimension of the state space, the width of the prior bounds, and the number of approximation nodes. We suggest iterating over the state space instead of solving the entire system each time. In each iteration step, we solve a simpler problem, which delivers conservative bounds on that state's policy and value function, which are updated and used in subsequent iterations. In this appendix, we show how to simplify Equation 6 and illustrate how to iterate to use newly found bounds on following iterations. We do not discuss, however, how to compute the piece-wise relaxation of the original constraints and refer the reader to the main text and Appendix A.1 for more details.

Let $\{\bar{V}_k(e), \underline{V}_k(e), \bar{P}_k(e), \underline{P}_k(e)\}_{e\in\{1,..,M\}^2}$ be known bounds of the price and value function at state e at iteration k. These are usually loose bounds, which ensure that all relevant equilibria lay within them. Our algorithm iterates over the state space a fixed number of times N_{inner} (according to a specific rule that will be defined below) and updates the price and value bounds for those specific states. These newly updated bounds are then used to update the bounds in subsequently visited states. After updating the bounds N_{inner} times—where some states can be visited more frequently than others—we call the resulting bounds $\{\bar{V}_{k+1}(e), \underline{V}_{k+1}(e), \bar{P}_{k+1}(e), \underline{P}_{k+1}(e)\}_{e\in\{1,..,M\}^2}$, and compare them to bounds at k, to establish our stopping rule. We continue iterating on bounds until the difference between $k+1$ and k bounds is less than a setup criterion or a maximum number of iterations N_{outer} is reached.

The algorithm consists of three key steps: 1) how to update bounds at state e, 2) how to transition from state e to e', and 3) a stopping rule. Next, we describe each step.

Updating bounds at state e To find the best possible bounds, we need to include $M^2 \times 2$ relaxed equilibrium conditions. Including all these constraints can be computationally intensive due to the number of discrete variables involved. However, some equilibrium conditions contain more information than others. Consider the case of state (M, M), the Bellman and FOC are

$$\begin{aligned} &\text{[Bellman]} \quad V^*(M,M) = D_1^*(M,M)(p^*(M,M) - c(M)) + \beta D_1^*(M,M)\bar{V}_1(M,M) + \beta D_2^*(M,M)\bar{V}_2(M,M) \\ &\text{[FOC]} \quad \sigma = D_1^*(M,M)(p^*(M,M) - c(M)) - \beta\bar{V}_1(M,M) + \beta D_1^*(M,M)\bar{V}_1(M,M) + \beta D_2^*(M,M)\bar{V}_2(M,M) \end{aligned} \tag{7}$$

where, $\bar{V}_1$ and $\bar{V}_2$ are known functions of the values over contiguous states

$$\begin{aligned} \bar{V}_1(M,M) &= (1-\delta)^M V^*(M,M) + \left\{1-(1-\delta)^M\right\} V^*(M,M-1)) \\ \bar{V}_2(M,M) &= (1-\delta)^M V^*(M,M) + \left\{1-(1-\delta)^M\right\} V^*(M-1,M)) \end{aligned}$$

When forgetting values are low, Equation 7 is ”almost” a system of two equations on two unknowns:

$V^*(M,M), p^*(M,M)$.[54] To fully determine the value of $V^*(M,M), p^*(M,M)$ we only need information on $V^*(M-1,M)$ and $V^*(M,M-1)$. Furthermore, loose bounds on these values can still be quite informative because its influence on $V^*(M,M), p^*(M,M)$ is scaled down by the large forgetting parameter. Therefore, taking loose bounds of $V^*(M-1,M)$ and $V^*(M,M-1)$, $\bar{V}(M-1,M), \underline{V}(M,M-1)$, the piece-wise linear bounds of equilibrium conditions in Equation 7, should still provide valuable information about bounds at (M,M).

More generally, the influence of state (s_1,s_2)'s value on the state (e_1,e_2)'s equilibrium outcomes is determined by the probability of reaching state (s_1,s_2) from (e_1,e_2) in subsequent plays. Furthermore, this probability is determined by equilibrium prices. Therefore, we use our price bounds to determine which equilibrium conditions influence the equilibrium at state (e_1,e_2) the most.

Let $\Omega^{2*}_{e\to s}$ denote the probability of transitioning from the state e to state s in two periods at equilibrium. Equilibrium prices determine $\Omega^{2*}_{e\to s}$ through the equilibrium demand one step ahead and two steps ahead. Thus, we can also define the two-step transition for any pair of such demands $\Omega^{2}_{e\to s}(D_1,D_2)$, where it is understood that D_1 and D_2 are functions from M^2 to $[0,1]$.

Hence, to update bounds at state e, we simplify Equation 6 to include piece-wise linear envelopes of equilibrium constraints[55] for all state s such that satisfies $\Omega^{2}_{e\to s}(D_1,D_2) > \gamma$, where $\gamma \in [0,1]$ is a tuning parameter to be chosen. It is sometimes useful to make this parameter a function of the iteration $it_{outer} \in \{1, ..., N_{outer}\}$ and the size of the update between k and $k+1$ iterations. As a general rule, we want to include more equilibrium constraints further along in the iteration or when our current constraints are not making enough progress in bounds updates.

Let $H(e)$ be the equilibrium outcome at e we are trying to bound. Then, instead of solving the large optimization problem in Equation 6, we solve the much simpler problem,

$$\begin{aligned}
\bar{H}(e), \underline{H}(e) = \max_{\mathbf{V}^*,\mathbf{p}^*} \text{ and } \min_{\mathbf{V}^*,\mathbf{p}^*} \quad & H(e) \\
\text{s.t.} \quad & \underline{\mathbf{F}}_s(\mathbf{V}^*,\mathbf{p}^*) \le 0 \le \overline{\mathbf{F}}_s(\mathbf{V}^*,\mathbf{p}^*) \quad \text{for all } s \in S(e) \\
& \underline{\mathbf{V}} \le \mathbf{V}^* \le \overline{\mathbf{V}} \\
& \underline{\mathbf{P}} \le \mathbf{P}^* \le \overline{\mathbf{P}}
\end{aligned} \tag{8}$$

where $S(e) = \{s \in M^2 : \Omega^{2}_{e\to s}(D_1,D_2) > \gamma\}$

Although we only include piece-wise linear bounds for a subset of equilibrium constraints, Equation 8 uses all known bounds on prices and values -indeed, we can also include all constraints that enter linearly into the problem.

Finally, observe that (D_1,D_2) are undetermined since we do not know equilibrium prices. We use our previously computed price bounds to compute demand bounds, exploiting the symmetry of the game. That is, for any state $s = (s_1,s_2)$

[54]This statement is exactly right when there is no forgetting, and the problem can be solved by backward induction.

[55]$\{\underline{F}^1_s, \underline{F}^2_s, \bar{F}^1_s, \bar{F}^2_s\}$

$$\frac{1}{(1+exp(\frac{\overline{P}(s_1,s_2)-\underline{P}(s_2,s_1)}{\sigma}))} = \underline{D}(\underline{P},\overline{P})(s_1,s_2) \leq D(s_1,s_2) \leq \overline{D}(\underline{P},\overline{P})(s_1,s_2) = \frac{1}{(1+exp(\frac{\underline{P}(s_1,s_2)-\overline{P}(s_2,s_1)}{\sigma}))} \tag{9}$$

Each time we visit a state e, we choose D_1 and D_2 as a convex combination of the lower and upper bounds $\underline{D}(\underline{P},\overline{P}), \overline{D}(\underline{P},\overline{P})$, whose weights are drawn from a standard uniform distribution.

State transition By solving Equation 8, we obtain new bounds on prices and values at e. Then, we must choose how to move into the new state e'. There are several options for choosing how to transition from one state to the next, and the optimal approach depends on the application at hand. The most natural one is to use a modified one-step transition kernel $\Sigma_{e\to e'} = \tilde{\Omega}_{e\to e'}(D_1)$, such that $\tilde{\Omega}_{e\to e}(D_1) = 0$. $\tilde{\Omega}_{e\to e'}(D_1) = \Omega_{e\to e'}(D_1) \times C \quad \forall e' \neq e$.[56] Although this is an attractive option, bounds of states that are unlikely to occur, given the forgetting parameter, will be visited too infrequently and remain too wide. In a way, this procedure parallels the stochastic algorithm in Pakes and McGuire (2001). So, we adjust the actual transition probabilities over-weighting states that the algorithm has visited infrequently.

We explore different ways to implement this over-weighting. First, if the equilibrium is unique, the actual transition kernel of the game forces the algorithm to remain within the recurrent class –if bounds on prices were sharp, we could again establish a parallel with Pakes and McGuire (2001). In this case, we would like to overweight states visited infrequently but visited at least once by the algorithm in previous steps. In that way, we would remain within the recurrent class –if we started from an element in the recurrent class– but states would be visited more evenly. With multiple equilibria, states in one recurrent class could communicate with states in other recurrent classes even when bounds are sharp due to equilibrium selection. In this case, we would prefer to update bounds over the whole state space. Therefore, the over-weighting of infrequently visited states should also include those not visited by the algorithm until that iteration.

Stopping Rule Finally, we establish a criterion to stop the algorithm. We determine the number of iterations between criterion updating, N_{inner}, and define the criterion as the update on the value function's bound during the last N_{inner} iterations.

The whole procedure is summarized by Algorithm 2:

[56] Here, C is a constant to ensure transition probabilities add up to one.

ALGORITHM 2

Iterative stochastic algorithm to solve equilibrium bounds of a dynamic game.

Choose e_0 and γ_0, set $Freq[e] = 0 \quad \forall e$
while $crit > tol$ **do**
  i = 0
  for $i \leq N_{inner}$ **do**
    Bounds updating

$D_{i,k}[e_{i,k}] \leftarrow \underline{P}_{i,k}[e_{i,k}], \overline{P}_{i,k}[e_{i,k}]$

$\Omega^2_{i,k} \leftarrow D_k, \gamma_k$

$\underline{P}_{i+1,k}[e_{i,k}], \overline{P}_{i+1,k}[e_{i,k}], \underline{V}_{i+1,k}[e_{i,k}], \overline{V}_{i+1,k}[e_{i,k}] \leftarrow_{\text{solve Equation 8}}$

State Transition

$D_{i+1,k}[e_{i,k}] \leftarrow \underline{P}_{i+1,k}[e_{i,k}], \overline{P}_{i+1,k}[e_{i,k}]$

$\Sigma_{e_{i,k} \to e'} \leftarrow D_{i+1,k}[e_{i,k}], Freq_{i,k}$

$e_{i+1,k} \leftarrow$ drawn from $\Sigma_{e_{i,k} \to e'}$

$Freq[e_{i+1,k}] = Freq[e_{i+1,k}] + 1$
  end for
  Stopping Rule

---

crit = $||\underline{V}_{k+1} - \underline{V}_k|| + ||\overline{V}_{k+1} - \underline{V}_k||$ with $k+1$ representing the same iteration as N_{inner}, k
$k + 1 \leftarrow k$
end while

B Supplementary material for numerical examples.

B.1 Additional material for discrete mixed logit example

B.1.1 Profit Function and Best Responses

Figure B.1: Firm's profit under several rival prices

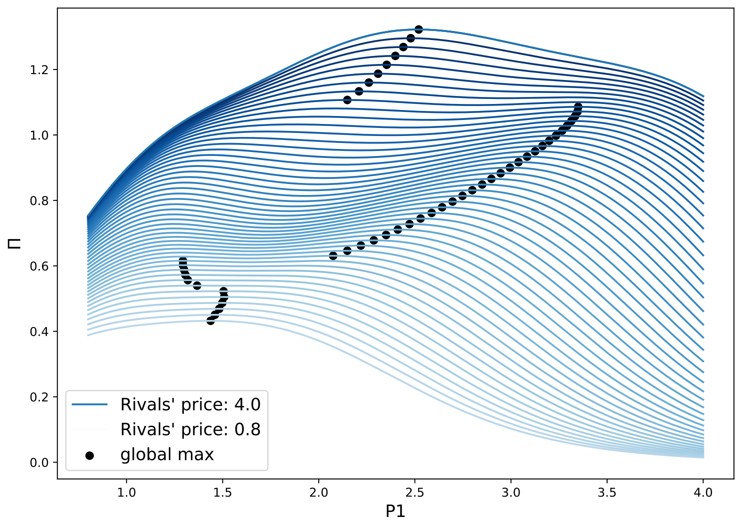

Notes: . Shopper sensitivity = -3.7.

Figure B.2: Firm's Best Responses under different shoppers' price sensitivities

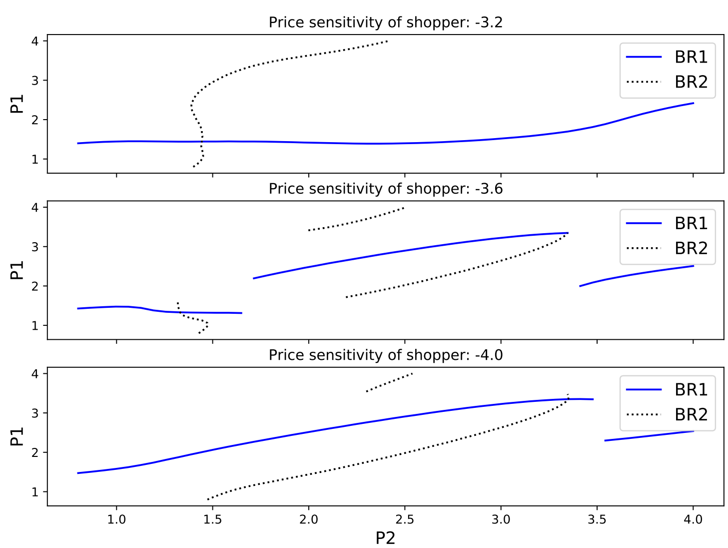

B.1.2 Non-linear solutions

Figure B.3: Non-linear solution of Equation 3

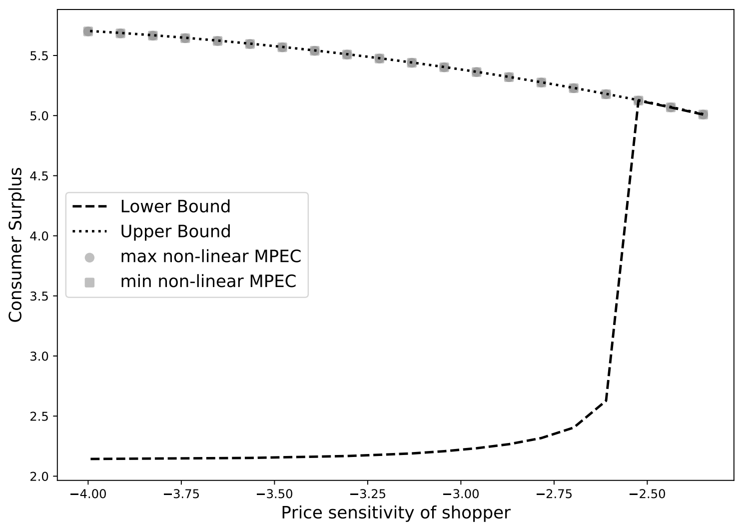

B.2 Additional material for discrete choice with conditional heteroskedasticity.

B.2.1 Non-linear solutions

Figure B.4: Equilibrium prices with conditional heteroskedastic logit demand, found with non-linear solvers.

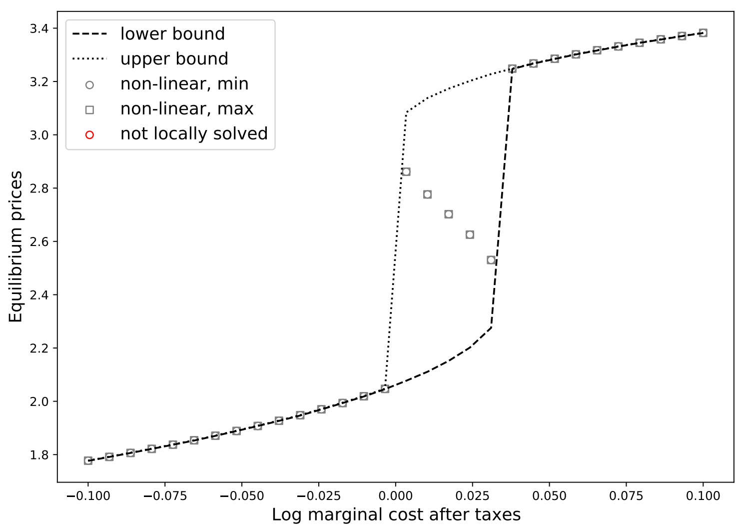

Notes: We solve the system of equations using Knitro within JuMP, an optimization environment written for Julia. We allow Knitro to initialize multiple start values -ms_enable is set to 1. Black lines -dashed and dotted- indicate conservative bounds on equilibrium prices.

B.3 Additional material for dynamic game example

B.3.1 Non-linear solutions

The equilibria of the game can be computed by directly solving the non-linear system of equations.

$$\mathbf{F}(\mathbf{V}^*, \mathbf{p}^*) = \begin{bmatrix} F^1_{(1,1)}(\mathbf{V}^*, \mathbf{p}^*) \\ F^1_{(2,1)}(\mathbf{V}^*, \mathbf{p}^*) \\ \vdots \\ F^2_{(M,M)}(\mathbf{V}^*, \mathbf{p}^*) \end{bmatrix} = 0$$

We solve the non-linear system of equations using Knitro. We identified two equilibria of the game at $\delta = 0.1218$. Figure B.5 presents these equilibria's policy and value function.

Figure B.5: Two known equilibria of the game at $\delta = 0.1218$.

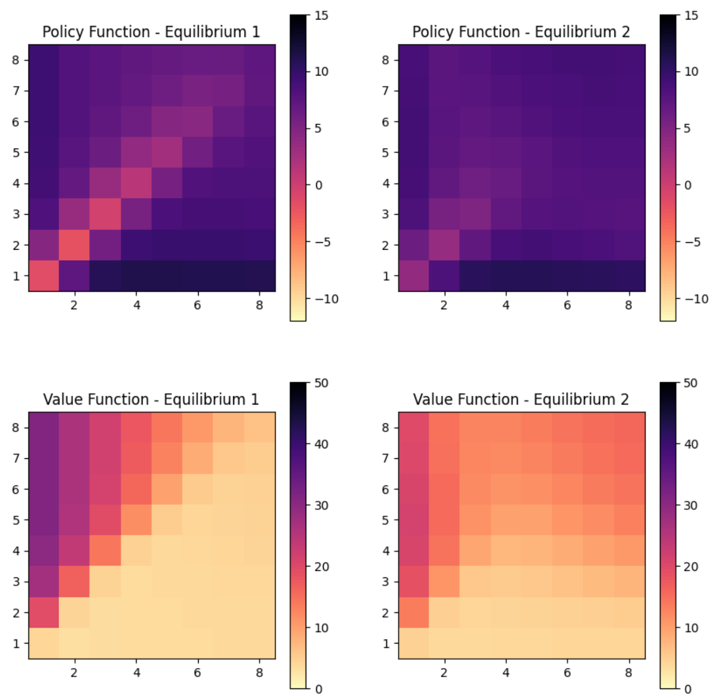

Notes: We solve the system of equations using Knitro within JuMP, an optimization environment written for Julia. We are allowing Knitro to initialize multiple start values -ms_enable is set to 1.

B.3.2 Other Figures

Figure B.6: Approximation of the function $\frac{\sigma}{1-D(z)}$

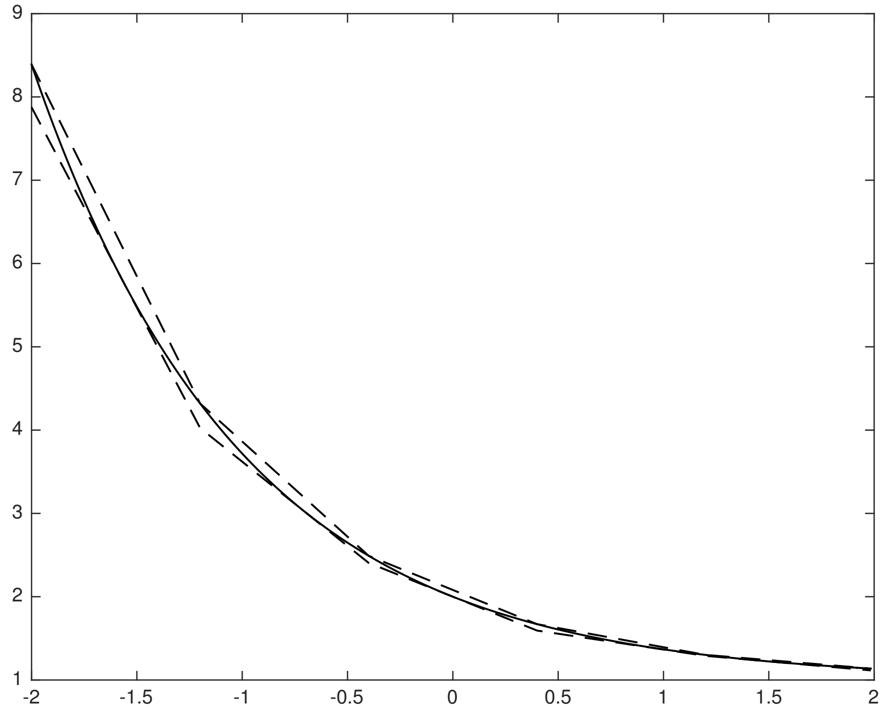

Notes: Approximation made for $J = 5$ and $z \in [-2, 2]$..

Figure B.7: Bounds on equilibrium prices and values state by state. $\delta = 0$