Bounding Outcomes in Counterfactual Analysis*

Mar Reguant
Northwestern University and NBER

First version: January 15, 2015
This version: August 25, 2016

PRELIMINARY AND INCOMPLETE
Click here for latest version

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 difficult to exhaustively characterize. I propose a new methodology 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, I propose to reframe equilibrium constraints as a relaxed mixed-integer linear program. I show that the framework can also be used to narrow down equilibria, by imposing additional equilibrium constraints. I provide examples related to multi-unit auctions and dynamic games, two areas in which counterfactual analysis has proven elusive.

*Email: mar.reguant@northwestern.edu. I 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 difficult to exhaustively characterize. I propose a new methodology 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 (and possibly more). Because the bounds are conservative, they are easier to compute than tight bounds, albeit being less informative.

The motivation for the method is based on common challenges that are 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 that predict the impact of policies out of sample. To do so, the researcher poses first a model, which is estimated, and finally 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 meaningfully capture the observed patterns in the data. At the same time, the solution to the posed sufficiently-rich 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 in instances in which 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 that we can consider, e.g., to mechanisms that are truthful (Kastl, 2011; Hortaçsu and McAdams, 2010; Kang and Puller, 2008).

I present a new method to bound counterfactual outcomes under relatively general conditions. The basic idea of the method is intuitive. Imagine that 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 when it comes to estimation.

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 (or at least known to belong to a set). I 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 constraints (it could also be inequalities) that describes the equilibrium conditions of the game. Note that the method only searches for the minimum and maximum outcome that can be achieved when equilibrium conditions are satisfied. This is in contrast with homotopy methods, which try to characterize as many equilibria as possible, but has the advantage of being more computationally tractable.

To perform counterfactual experiments, the researcher can compute bounds on a given outcome 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 they will be dependent on which equilibrium is being played.

To fix ideas, consider a simple static entry game in which multiple equilibria exist.[3] The example is inspired by Fowlie et al. (2014), who consider entry and exit decisions by cement plants. Suppose that there is a dirty firm with large capacity and a clean firm with smaller capacity, which are competing in quantities. To stay in the market, they have to pay a fixed cost. 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, depending on the carbon tax, multiple equilibria might arise. In particular, for medium ranges of carbon taxes (between \$13 and \$24, approximately), either the dirty or the clean plant might exit the market, leading to substantially different levels of 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. One can see that, in such case, the outcomes under the carbon tax plus a subsidy are more tilted towards the lower emissions equilibria. In particular, a carbon tax of \$25 dollars plus a \$10 output subsidy

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

[3]The presence of multiple equilibria is common in entry/exit game of complete information (Bresnahan and Reiss, 1990).

Figure 1: A simple example of counterfactual bounds

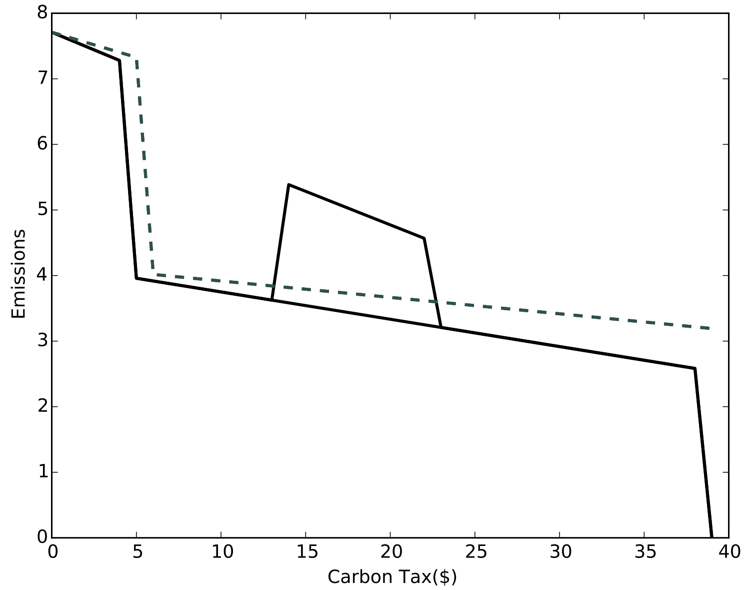

A simple entry game in which there is a dirty firm with quadratic costs $C(Q_1) = 20Q_1 + Q_1^2 + 1.5\theta$, and a clean firm with quadratic costs $C(Q_2) = 20Q_2 + 2Q_2^2 + \theta$. Fixed costs of staying in the market are 100 and 80, respectively. Inverse demand is given by P = 120 - 10 Q. θ represents the carbon tax. The dashed line represents a situation in which a 40% output subsidy is introduced in the cap-and-trade scheme.

achieves the same equilibrium emissions as the lower bound when the carbon tax is $15, but does not support an equilibrium in which the dirty plant only stays.

This example is useful to convey that it is important to account for multiplicity of equilibria when computing counterfactual experiments. In this stylized setting, it is easy to check all equilibria in an exhaustive manner. One can also find the lowest and highest levels of a counterfactual outcomes more generally by finding the minimum and maximum emissions, subject to the participation constraints of the firms and their optimal quantity decisions being satisfied.[4] The advantage of searching for counterfactual bounds directly is that only those equilibria that give the most extreme equilibrium outcomes is 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, these counterfactual bounds might be good enough to allow us to make informed 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 problem are well behaved, then 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,

[4]See appendix for a full characterization of this simple game.

but not necessarily all of them.[5]

I 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.[6] This presents several advantages. First, the linearity, together with the use of integer variables, ensures that the solution can be guaranteed to be a global optimum. Second, because it is a relaxed version of the original problem, it necessarily includes the true bounds to the objective function. Therefore, even though the method relies on approximation techniques, it does not compromise the validity of the bounds. I 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. I 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, I show that the method can be used 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. 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 where marginal costs are infinity. I also show how to narrow down this relevant range in Proposition 4.

To demonstrate the applicability of the proposed methodology, I present two common applications in Industrial Organization in which equilibrium computation has been difficult. First, I show how the methodology can be used to compute bounds to equilibria in multi-unit auction, when firms compete in price-quantity schedules. This application is helps to highlight the advantages of the methodology, as to my knowledge, there were no previous tools that allowed for flexible counterfactual analysis in this setting. Second, I 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 has been previously identified as being susceptible to multiple equilibria using homotopy techniques (Besanko et al., 2010). I also use this example to highlight how the proposed methodology can be used as a powerful equilibrium refinement tool.

Related Literature. This paper is related to the literature examining how to compute equilibria in games, specially 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, I propose to bound the counterfactual outcome of interest in a conservative manner. The methods that I use are based 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).

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

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

Conceptually, the paper has links to the econometric literature proposing to estimate bounds for parameters of interest, instead of seeking for point identification, to make our statements more robust and less reliant on explicit assumptions (Manski, 1989, 1990, 2003).[7] Whereas the methods presented here focus on counterfactual computation, and take parameter estimates as given from a previous stage (point or set identified), the goal of the methodology is in the same spirit of ensuring that our statements, in this case about counterfactual outcomes, can be interpreted in a robust manner as worst case bounds.

There is also a growing literature on how to bound counterfactual outcomes in a robust manner.[8] Following on his work on partial identification, Manski (2007) shows how to perform robust counterfactuals in discrete choice settings with mild assumptions on preferences. Jia (2008) exploits supermodularity in an entry game to focus on highest profit equilibria, which enables the estimation and computation of the game. Uetake and Watanabe (2014) also study an entry game that has a lattice structure, and compute minimum and maximum counterfactual outcomes taking into account that the parameters of the model are only set identified. Sweeting (2009) takes into account 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) study an innovation game in which bounds to social surplus are computed from a set necessary equilibrium conditions that are easier to compute than the full equilibrium. I propose a framework that is potentially applicable to a wide range of problems. For the same reason, however, it might be less efficient than other alternative methods tailored to the specifics of a particular game or setting.

Finally, the examples presented are related to the literatures on supply function equilibria and dynamic games, which I discuss in the context of each particular application.

2 Methodology

I 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 represent 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 each of 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 it is true that prices could 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.

[7] There is a large and growing literature on partial identification and inference in such settings. See Tamer (2010) for a review.

[8] I focus here on papers that emphasize equilibrium and counterfactual computation. The current discussion is not intended to be exhaustive.

Assumption 2 limits the size of the strategy space to be finite. This could be relevant in continuous games, e.g., a game in which agents are choosing contingent plans over continuous states. To palliate this limitation, approximation techniques could be used 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. This means that all equilibrium constraints are considered at the same time. A one-step program has the advantage of avoiding iteration between the best responses of different agents, e.g., as typically done using a Gauss-Jacobi or Gauss-Seidel algorithm, where the best response of each firm is computed at a time given what others do at the current iteration. Whereas iteration methods can work well in practice, in many game settings one cannot guarantee that such procedure is a contraction mapping or that a unique equilibrium exists. For example, iteration methods might converge to particular regions of the strategy space. Solving all equilibrium constraints at the same time can be crucial to avoid the problems of "missing" particular equilibria in a game. The constraints of the program are the equilibrium constrains of the game.[9]

Importantly, the objective function is not necessary to define the equilibrium constraints of the game. Instead, the objective function is reserved to the counterfactual outcome that the researcher is interested in bounding. This is a key feature, as one can directly minimize and maximize the objective function to characterize the bounds.[10] 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 searchers for the maximum and minimum counterfactual outcomes that can be sustained in equilibrium.

To fix ideas, consider a setting in which the equilibrium constraints are necessary and sufficient, and can be expressed as a set of piece-wise linear constraints. In such case, the objective function can be minimized and maximized to obtain bounds on the counterfactual of interest.

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

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

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

[10]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).

and,

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

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. □

In this simplified environment, and given that the constraints are piece-wise linear, the bounds will be sharp, meaning that the end points are valid counterfactual outcomes satisfying the necessary and sufficient conditions to the game. This representation is therefore useful as a device to detect multiple equilibria. If there is a single equilibrium that satisfies $\mathbf{G}(\mathbf{x};\theta) = 0$, then $\underline{W}(\mathbf{x};\theta) = \overline{W}(\mathbf{x};\theta)$.[11] 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.[12] 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 particularly useful for the purposes of computing conservative bounds, that account for the computational burden of the program.

Even when the algorithm fails to exhaustively explore all potential solutions, branch-and-bound algorithms provide an upper and lower bound to the objective function (something called 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 provides 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 in which the objective function and the 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, I propose to approximate the equilibrium conditions of the game with piece-wise linear envelopes.[13] Linearity ensures that all solutions can be considered, even in the presence

[11] 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.

[12] 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.

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

of nonlinear constraints that would otherwise be difficult to explore. However, it comes at the expense of loosing 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 that define such bounds. Assume that we can define a counterfactual outcome that is piece-wise linear in equilibrium strategies, defined by* $W(\mathbf{x},\theta)$*.*

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* $[\underline{W}_u(\theta), \overline{W}_u(\theta)]$*. 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 special case of a relaxed optimization problem.[14] The optimum of a problem subject to relaxed constraints (easier to satisfy) 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. □

Proposition 2 generalizes the methodology to non-linear settings, at the expense of providing conservative bounds. Note that this setting can also accommodate the case in which the objective function is non-linear. 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.

The method provides valid bounds independently on how the pieces are defined (i.e., how the approximation breaks regions of the non-linear 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

[14] In integer programming, the "relaxed" term is often used to refer to a linearized version of a problem that convexifies integer variables. I 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.

which the pieces are considered is randomly generated.[15] 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, and indeed, can be used as additional refinements. This is important, as reducing the number of integer variables is an effective way to reduce 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 That said, the bounds might be very wide and uninformative in a given application.

Fortunately, an iterative approach can be used as a way to refine the approximation, while holding the number of approximation points fixed. Using the same bounding 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 narrow the range of approximation. Indeed, the range over which equilibrium constraints are approximated can be narrowed down successively, as suggested in Proposition 4. One key result is that such iterative search does not invalidate the conservative nature of the counterfactual bounds. This is in contrast 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* $\overline{W}_{uk}(\theta) \geq \overline{W}_{u,k+1}(\theta)$.

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

Figure 2: Successive Approximation of Equilibrium Constraints

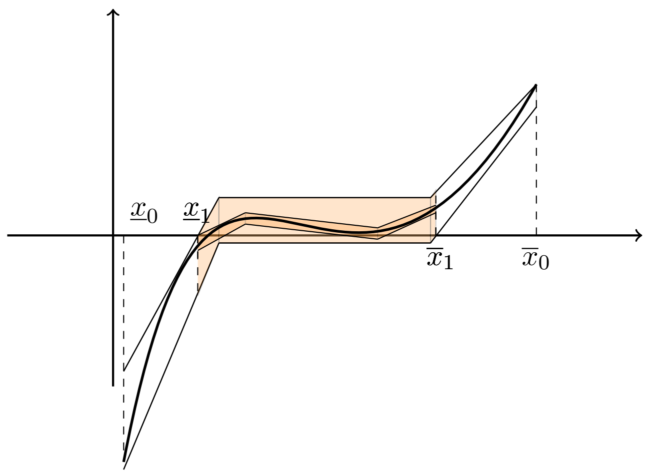

By narrowing down the relevant range of the action space, one can improve the approximation of equilibrium constraints, even when the number of piece-wise linear pieces is limited.

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. □

Figure 2 provides a graphical intuition for this procedure, in a single dimensional space for ease of depiction. In an intial iteration, the bounds on x are set to $[\underline{x}_0, \overline{x}_0]$ providing an initial envelope with three pieces. Under such envelop, one finds that at most $[\underline{x}_1, \overline{x}_1]$ can be sustained, narrowing down the range of approximation. With still three pieces, the envelopes are now much tighter.

Note that the procedure of maximizing and minimizing each strategic variable can also be useful to ensure that pre-specified limits in the range of a given variable x_i are not binding. As an additional practical matter, also note that depending on how costly step 3 is, one could decide to 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 be dependent on the problem at hand.

2.3 Relaxation

The proposed optimization approach is very amenable to alternative constraints, which can be useful in situations in which the equilibrium is difficult to solve for or situation in which not all equilibrium constraints can be characterized. Indeed, the approximation method described above is already a particular case of a "relaxed" problem, in which the solution to a non-linear game was transformed into a conservative set of piece-wise linear inequalities. 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, and therefore a feasible solution to the problem might be difficult to find in a reasonable amount of time. Indeed, mixed-integer programs are NP-hard, making them increasingly burdensome as the number of variables grow.

One kind of relaxed problems that is often considered in the optimization literature, is one in which integer variables are transformed into continuous one. 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 particularly useful in conjunction with Proposition 4. Consider a situation in which one is trying to narrow down the support of a particular strategic variable, e.g., in a Cournot game, one is trying to narrow down 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. I 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. Alternatively, one might not want to include all such constraints. A particular class of problems that is easy to fit in this framework is situations in which players are not exactly optimizaing, 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. As another example, it is quite common in the auction literature, that necessary local conditions are easier to characterize than full equilibrium conditions. To the extent that necessary conditions for optimality are informative, they can be used as a way to bound counterfactual outcomes. Because it is a relaxed problem, the bounds will still be valid, albeit larger. The informativeness of these bounds will be quite dependent on the particular application. I show how this method can be used in the context of multi-unit auctions, in Section 3.

Equilibrium Restrictions. Depending on the application, one might want to instead apply further restrictions to the equilibrium. For example, in particular economic applications, one might want to focus on equilibria that satisfy particular properties. Or maybe one would like to compare counterfactual bounds under alternative explicit equilibrium-selection rules. Another very natural restriction to add is second order conditions, to 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. I 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. If the solver completes an exhaustive search, and fails to find a set of equilibrium outcomes that are consistent with the relaxed equilibrium constraints, then it can provide numerical proof that an equilibrium does not exist that satisfies the original 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, on the contrary, the set of constraints were non-linear, we could only say that the solver did not find to manage a candidate equilibrium that satisfied the constraints, but we could not assert that no such equilibrium existed.

Accounting for Parameter Uncertainty. Parameter uncertainty can also be easily incorporated in this framework, i.e. uncertainty in θ. Imagine θ is known to belong to a set Θ, which has been obtained in the estimation stage.[16] In theory, one could solve for equilibrium bounds over several parameter combinations, sampling from the estimated distribution of parameters several times, and taking the minimum ever computed and the maximum ever computed. However, this can be very expensive and might miss equilibrium regions due to the finite nature of the sampling approach.

Fortunately, the proposed framework allows one to compute counterfactual bounds that are conservative and 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 Θ as the set in which 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 will 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. Therefore, the bounds might tend to be much more conservative than the implied coverage by the confidence interval or confidence set. 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.

[16] This set could either be obtained using explicit bound estimation, or for example as the set defined by the confidence intervals obtained in the estimation.

2.5 Discussion

The proposed approach presents some strengths and weaknesses that are important to discuss. The major 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 principles of non-iteration and piece-wise linear approximation, 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. A special case of relaxation is the case in which integer variables used for approximation are converted into continuous variable, 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 in which functions are highly non-linear, and which contain 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 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 to assess how it performs in practice, I present three examples from the literature in which limitations to counterfactual simulations have been previously identified. In Section 3, I consider a multi-unit auction. In Section 4, I consider a dynamic game. Finally, I discuss applications to discrete choice models.

3 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.[17] The inability to fully characterize supply function equilibria in a robust manner 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, in which firms choose a schedule of prices and quantities, which are 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 is also useful to show how the method can be applied in situations in which the strategy space is continuous, by structuring the necessary conditions of the game in a way that is amenable to the methodology.

3.1 Model outline

Consider a multi-unit auction game in which players can choose a schedule of prices and quantities. A typical example of such game are treasury bill auctions (Hortaçsu and McAdams, 2010; Kastl, 2011), or electricity markets (Wolak, 2003; Hortaçsu and Puller, 2008). At a given auction, there is an uncertain amount of units that are being auctioned, A_s. This amount is uncertain before the auction, and therefore is indexed with state $s = \{1, \ldots, S\}$.[18] Without loss of generality, I sort demand along states so that $A_s > A_{s-1}$.[19] Importantly, the dimenstionality of demand states is assumed to be finite.

To implement the methodology, I 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 the strategy of firms 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, in equilibrium the market clears, i.e.,

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

Figure 3 presents a depiction of 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, $q_{i,s+1} + q_{j,s+1} = A_{s+1}$.

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

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

[19] 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 3: Example of Candidate Supply Functions

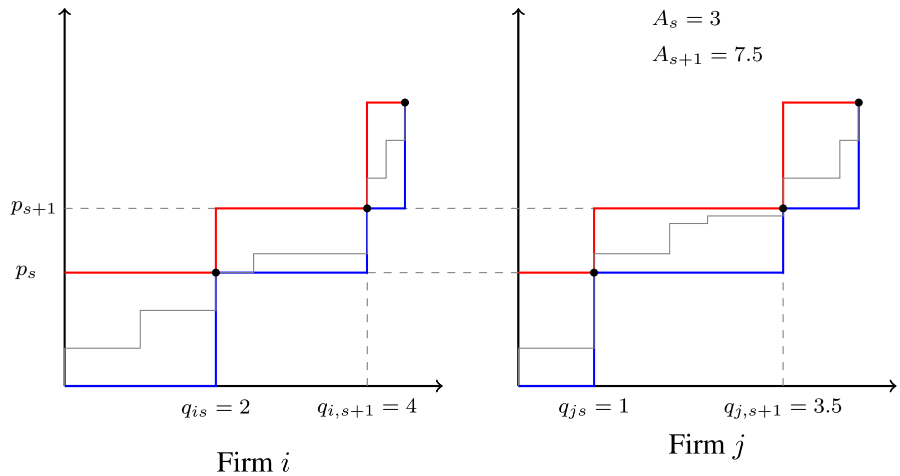

The pairs of offers $\{q_{is}, p_s\}$ and $\{q_{js}, p_s\}$ clear the market with demand $A_s = 3$ and price p_s. The pairs of offers $\{q_{i,s+1}, p_{s+1}\}$ and $\{q_{j,s+1}, p_{s+1}\}$ are clear the market when $A_{s+1} = 7.5$. The red envelope represents the upper bound to the supply function along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

Because I do not solve for the full schedule of offers, I 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 demand. 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?

3.2 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, as opposed to inferring the costs from the equilibrium strategies.

Because the schedule as defined is discrete, I consider deviations from optimal strategies taking a conservative version of other firms' strategies, i.e., taking the red and blue envelope as best and worst case scenarios to the supply function for other firms. Consider a situation in which firm i is offering 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, I 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}.[20] In such case, a lower bound deviation is given by,

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

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

where $C(.)$ is a weakly increasing cost function. The condition states that the firm should not be willing to increases 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 at least offer 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 that one does not need to force 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, to approximate revenues one can consider discretizing the price range and narrow down the envelopes at each state s iteratively. For the cost function, one can approximate this single-dimensional function that depends on quantity with piece-wise linear envelopes. In the special case of constant marginal costs, the cost function is already linear. These constraints could be expanded, for example to consider deviations of several steps at the same time.

3.3 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.[21] To do so, one can solve the 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.

[21] 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).

Figure 4: Bounds on Procurement Outcomes

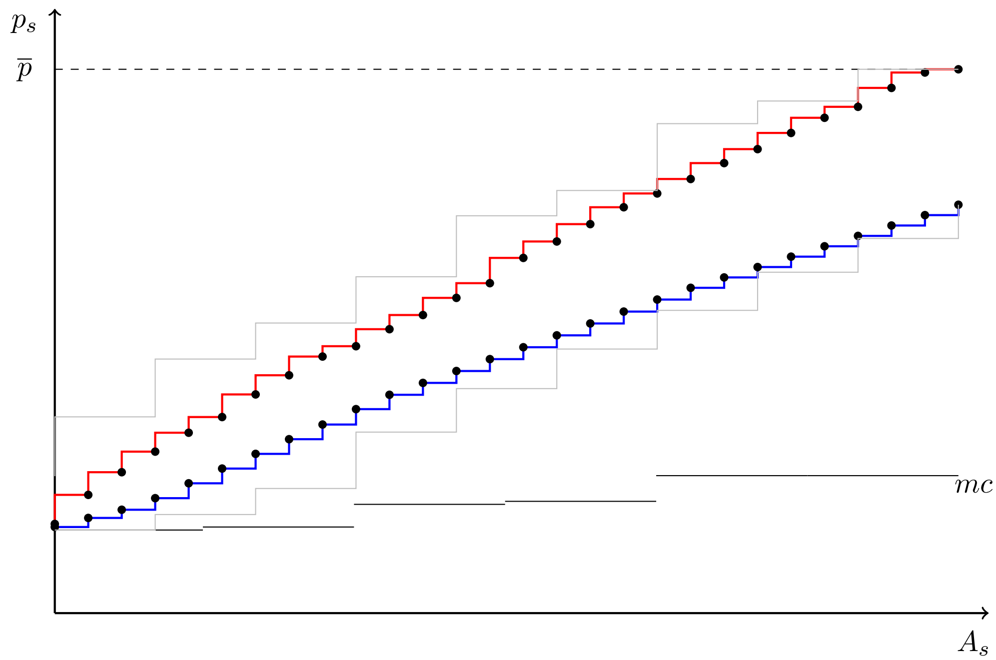

The red bound shows the supply curve that generates the highest procurement costs, whereas the blue curve shows the supply curve that generates the lowest procurement costs. These supply curves are consistent with necessary equilibrium conditions, and provide valid conservative bounds to procurement 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.

Example 1. I consider an example related to electricity auctions. There are two firms who are competing in a multi-unit auction. Each of them 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 constraints the maximum price that firms can set. I solve for the optimal strategies of the firms given a demand distribution $\{A_1, \ldots, A_S\}$. Uncertainty in demand is coming from residual uncertainty about other firms and/or renewable production, rather than demand in itself, which is typically forecasted with a high degree of accuracy. I assume that the two firms have substantial market power, and therefore might become pivotal, i.e., for some state s their production is needed to cover demand. I also assume that the lowest demand realization is above zero, i.e. $A_s > 0$, which in the context of electricity markets makes sense, as uncertainty is usually bounded.

A typical concern to policy makers in electricity markets is the expected cost of electricity to consumers, i.e., $\sum_s \omega_s p_s A_s$. Therefore, I look for bounds on procurement costs of electricity, allowing the equilibrium concept to be quite flexible (i.e., only imposing the equilibrium constraints described above). Figure 4 shows the aggregate supply curves that generate the highest (red) and lowest (blue) procurement costs. One can see that even the scenario with lowest procurement costs features a supply curve above marginal costs. The scenario with highest procurement costs is below the price cap, except for the highest demand realization.

3.4 Discriminatory vs. Uniform Auction

TBA

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 a variety of industries (Benkard, 2004; Jofre-Bonet and Pesendorfer, 2003; Ryan, 2012). Whereas dynamic interactions are relevant in strategic settings, implementation of models of industry dynamics have been limited. Two factors are important in explaining this limited implementation are the computational burden of computing Markov equilibria, as well as our limited capability of exhaustively characterizing such equilibria.

To address the first problem, researchers have proposed relaxing the equilibrium concept (Weintraub et al., 2008; Abbring et al., 2016) or relying on approximating techniques (Farias et al., 2012). The tools presented here aim to address the second problem, and can complement other techniques such as homotopy methods. They can also be applied to alternative equilibrium concepts that are less computationally demanding, such as oblivious equilibria, to the extent that multiple equilibria is still a concern in such setting.[22] The tools can also be used to partially characterize equilibria in dynamic games in a way that is less computationally intensive, 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.[23] The game has two firms $n \in \{1, 2\}$, each with a current state $e_n \in \{1, ..., M\}$. This state reflects their cost function: the higher the states, the lower the firm's marginal cost is. In particular, he 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:

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

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

[22] 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.

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

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 case, 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 looses some of the accumulated stock and suffers an increase in marginal cost, i.e., a reduction in their state.

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

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

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

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

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,

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

Reassuringly, the problem satisfies conditions for existence of an equilibrium (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, one can find bounds to the solutions to the system of equations by means of approximation techniques. As explained above, the system of non-linear equations is substituted by conservative piece-wise linear constraints.

I propose to find bounds to particular outcomes of interest. For example, in the context of learning-and-forgetting games such as the one above, which has a race feature, it might be interesting to assess the minimum and maximum net present value at time zero, when both firms have high marginal costs and have not learned yet, i.e., when the state is equal to one for both firms.[24] In such case, one can solve two separate programs,

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

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 true equilibrium conditions $\mathbf{F}(\mathbf{V}^*,\mathbf{p}^*)$.

It is important to note that when solving for the maximum and minimum $V(0,0)$, the program is still solving for all variables, $\{\mathbf{V}^*,\mathbf{p}^*\}$. However, the method is focusing on finding the upper and lower bound to the value function at that particular state, while finding values for all the other variables that are consistent with the relaxed equilibrium constraints.

4.2 Envelopes on FOC

Similar to the discrete choice case, it is useful to show how the approximation techniques work with the simpler condition, which in this case is the first-order condition. The FOCs at each state $\mathbf{e} = \{e_1, e_2\}$, are given by,

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

One can see that most of the price and value function terms enter the conditions in a linear fashion. The only non-linear part is given by $p_1 - p_2$, which enters the demand function.

One can transform these equations into necessary 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

[24] One can also compute bounds to the all prices and value functions.

Figure 5: Approximation of the function $\frac{\sigma}{1-D(z)}$

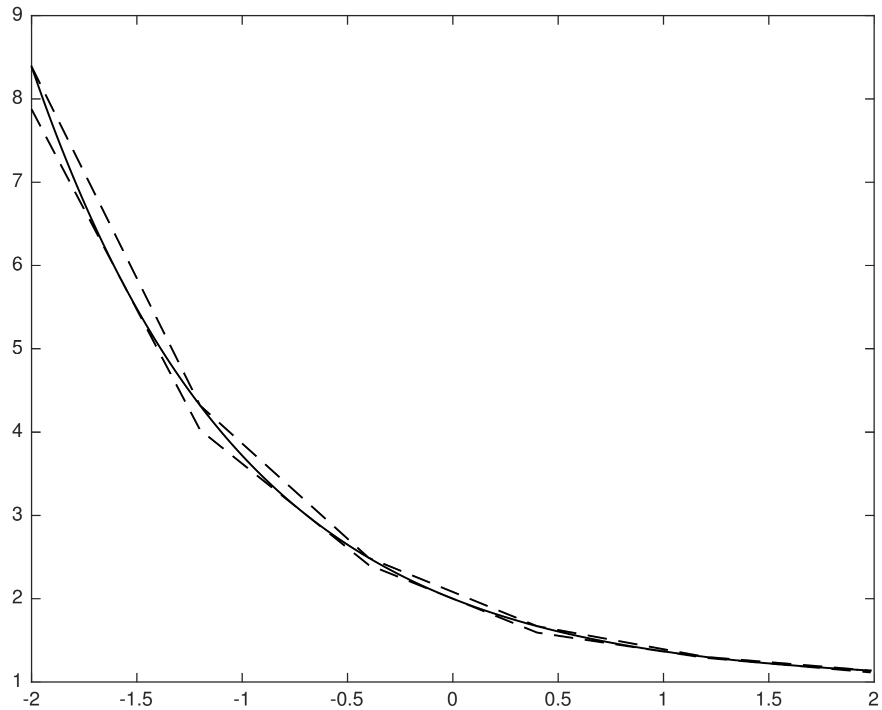

$\{p_1, p_2, z\}$,

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

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.

Example 2. Consider the case in which $\sigma = 1, \rho = 0.85, \kappa = 10$. For $\beta = 0$, one can solve for the equilibrium bounds at each state $\mathbf{e}$. Consider approximating the function $f(z) = \frac{\sigma}{1-D(z)}$ with a piece-wise linear approximation with J equally pieces between $z \in [\underline{z}, \overline{z}]$. As $J \to \infty$, the upper and the lower bound converge to the unique equilibrium of the game. Figure 5 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 bounds to the function are relatively accurate.

Table 1: 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 $[\underline{p}_1,\overline{p}_1]$	[10.07,10.13]	[10.03,10.19]	[10.08,10.10]	[10.07,10.11]
Actual Average Prices p_1	10.09	10.09	10.09	10.09
Average $(\overline{p}_1-\underline{p}_1)$	0.07	0.16	0.02	0.04
Maximum $(\overline{p}_1-\underline{p}_1)$	0.10	0.23	0.03	0.06
Average $(\overline{p}_1-\underline{p}_1)/\underline{p}_1$	0.01	0.02	0.00	0.00
Maximum $(\overline{p}_1-\underline{p}_1)/\underline{p}_1$	0.01	0.02	0.00	0.01

Notes: Results from Example 2 with $m=5$, $M=8$, $\beta=0$, $\sigma=1, \rho=0.85, \kappa=10$. The table presents averages across M^2 states. Note that in the static case, it is easy to characterize the actual unique equilibrium prices at each state, which are also presented in the table as a point of reference.

With this approximation of f, one can define bounds on the equilibrium conditions that are piece-wise linear, and compute counterfactual bounds in a robust fashion. In this case, I compute bounds on equilibrium prices. Table 1 shows the accuracy of the method in this simplified example. Because the example is simple, one can compare the counterfactual bounds to the actual unique equilibrium. One can see that, even with relatively coarse approximations, the bounds provided are relatively tight. In all four cases considered, there is at most a 1% to 2% difference between the upper and the lower bound.

4.3 Envelopes on Bellman

To solve for the full equilibrium, one requires the Bellman equations as well. Making use of the fact that $D_1(z)+D_2(-z)$ = 1, the Bellman conditions become, [25]

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

In the Bellman equations, the non-linear terms affect both prices and value functions. In particular, $D_n(z)$ appears interacted 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 considered, or a linear combination across different value functions (i.e., $\overline{V}_{n2}$).

The function $g(z,V)$ needs to be approximated as a function of both z and V. There are several ways to approximate this function. One option is to use interaction terms between z and V and create a lower and upper envelope to $g(z,V)$ by means of planes. An alternative approach is exploit the fact that V appears multiplicative to z, similar to the case of product prices in the BLP example. The particular approximation choice might also be affected by the computational intensity of the method.[26]

[25] See appendix in Besanko et al. (2010) for a derivation. Note that I use a slightly different notation.
[26] See details in appendix. TO BE ADDED

Table 2: Counterfactual bounds on prices for the dynamic case

	$\delta = 0.0$	$\delta = 0.0275$	$\delta = 0.2$
Average bounds $[\underline{p}_1, \overline{p}_1]$	[8.92,8.93]	[0.87,13.59]	[6.20,6.22]
Iterated Prices p_1	8.92	8.87	6.21
Average $(\overline{p}_1 - \underline{p}_1)$	0.01	12.71	0.01
Maximum $(\overline{p}_1 - \underline{p}_1)$	0.10	23.87	0.06
Average bounds $[\underline{V}_1, \overline{V}_1]$	[19.92,19.93]	[6.18,35.88]	[15.02,15.08]
Iterated Values V_1	19.92	19.55	15.05
Average $(\overline{V}_1 - \underline{V}_1)$	0.02	29.70	0.06
Maximum $(\overline{V}_1 - \underline{V}_1)$	0.09	38.01	0.13

Notes: Results from Example 2 with $m = 5$, $M = 8$, $\beta = 1/1.05$, $\sigma = 1, \rho = 0.85, \kappa = 10$. The table presents averages across M^2 states. The solution to the game using a value-function iteration approach are also reported.

Example 2 (continued). Following upon the previous example, consider now the case in which $\beta = 1/1.05$ and $\delta = \{0.0, 0.0275, 0.1\}$. As shown in Besanko et al. (2010), the game can lead to multiple equilibria, specially in the range where $\delta = 0.0275$. The bounds approach finds that the bounds are perfectly tight for $\delta = 0$, which can be shown to feature a unique equilibrium. For $\delta = 0.0275$, the method finds much larger bounds. It is indeed a range in which multiple equilibria arise, some that are very aggressive and some that resemble the equilibrium without forgetting, in which prices are mostly flat. In the table, only average bounds are reported. Checking the bounds for each state confirms that the kind of equilibria encountered using the homotopy methods are contained in the bounds. Finally, when forgetting is sufficiently high and $\delta = 0.2$, the game has again a unique equilibrium, and the bounds become tight. Indeed, the absence of multiple equilibria as δ grows was also found using the homotopy methods in Besanko et al. (2010).

The table also provides a comparison to the equilibria using a value-function iteration approach. One can see that the bounds successfully bound the equilibrium that is also encountered by iteration.[27] When $\delta = 0$, we know there is a unique equilibrium, which coincides with the iteration approach. Importantly, when $\delta = 0.0275$, the iteration approach picks up the more stable equilibrium in which prices are mostly flat. The bounds approach, on the contrary, detects a potential for multiple equilibria. Finally, the results also coincide when $\delta = 0.2$. Using the bounds method we not only confirm that the two approaches coincide, but also that the equilibrium is indeed unique.

4.4 Equilibrium Selection

Counterfactual bounds that are very wide, as reported in Table 2, can be quite uninformative. Often, the researcher would like to put some additional constraints to 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 the features of the game or the institutional environment.

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

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

One possibility is to restrict equilibria to situations in which, *ceteris paribus*, firms are better off when their marginal costs are lower.[28] Using the proposed methodology, this can be easily embedded into the program by adding the following constraints:

$$V_1(e_1, e_2) \geq V_1(e_1', e_2) \text{ if } C(e_1) < C(e_1'),$$
$$V_2(e_1, e_2) \geq V_2(e_1, e_2') \text{ if } C(e_2) < C(e_2').$$

Failure to solve for the equilibrium conditions of the game under these conditions implies that there are no equilibria that satisfy such rule, which can when firms compete very aggressively in equilibrium, and no "flat equilibria" exist.

There are many other alternative ways to narrow down 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.

Besanko et al. (2010) classify equilibria according to pricing behavior, by examining whether firms price below marginal cost at some states, which highlights the dynamic incentives to learn in the game. Using the methodology proposed in this paper, one could could rule out such equilibria by simply imposing that,

$$P_1(e_1, e_2) \geq C_1(e_1),$$
$$P_2(e_1, e_2) \geq C_2(e_2).$$

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.,

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

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

Table 3: Counterfactual bounds on prices for the dynamic case

	No restrictions	Monotonicity V in Own Costs	Price above Own Costs
Average bounds $[\underline{p}_1, \overline{p}_1]$	[0.87,13.59]	[7.61,9.90]	N/A
Iterated Prices p_1	8.87	8.87	8.87
Average $(\overline{p}_1 - \underline{p}_1)$	12.71	2.29	
Maximum $(\overline{p}_1 - \underline{p}_1)$	23.87	14.45	
Average bounds $[\underline{V}_1, \overline{V}_1]$	[6.18,35.88]	[16.75,22.72]	N/A
Iterated Values V_1	19.55	19.55	19.55
Average $(\overline{V}_1 - \underline{V}_1)$	29.70	5.93	
Maximum $(\overline{V}_1 - \underline{V}_1)$	38.01	18.63	

Notes: Results from Example 2 with $m = 5$, $M = 8$, $\beta = 1/1.05$, $\sigma = 1, \rho = 0.85, \kappa = 10$, $\delta = 0.0275$. The table presents averages across M^2 states. The solution to the game using a value-function iteration approach are also reported.

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

Example 2 (continued). Table 3 presents the results from the dynamic game when $\delta = 0.0275$. Column one presents the same bounds as in Table 2. Column two, on the contrary, restricts equilibria to be monotonic in own marginal cost. As one can see, the restriction appears to limit some of the equilibria. It is interesting to note that the equilibrium picked up by the iteration method is still within the bounds, presumably because it is the one that features smoother policy and value functions.Considering constraints on prices so that firms price above marginal cost, the algorithm fails to find an equilibrium, implying that no equilibrium exists in which firms do not price below their marginal cost at some state.

5 Other Applications

The method that I have presented can be applied to an array of environments and games. I present here two more examples, a Nash-Bertrand pricing game with heterogeneous consumers, and a Nash-Bertrand pricing game with attribute choice. Another example where the proposed methodology can be useful is capacity-constrained Cournot models. Indeed, Ito and Reguant (2016) use the proposed methodology to check for multiple equilibria in a model of sequential electricity markets.

5.1 Discrete Choice with Consumer Types

The framework that I have presented can also be used in discrete choice models with random coefficients, which have been widely used to perform counterfactual analysis.[29] Discrete choice models with endogenous prices might present multiple equilibria in the presence of heterogeneous consumers and/or multi-product

[29] See Berry et al. (1995, 1999); Nevo (2000b); Petrin (2002); Leslie (2004), among others.

firms (Berry et al., 1995; Echenique and Komunjer, 2009).[30] The presence of multiple equilibria can pose issues for the estimation of these demand systems. Dubé et al. (2012) show how an MPEC approach can be used to estimate the likelihood function, while remaining agnostic about which equilibrium is being played in the data. Even conditional on having consistent estimates of the demand and supply parameters, counterfactual simulations are subject to the same multiplicity concerns. This paper presents methods to address the later challenge, taking the value for the fundamentals as given.

Consider a market with $j = 0, \ldots, J$ products, including the outside option indexed by zero. There are $n = 1, \ldots, N$ firms, each of whom sells 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 that is approximated with $i = 1, \ldots, ns$ consumer types, with individual product-specific tastes μ_{ij}.

In order to solve the pricing equilibrium, one needs to solve for 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,

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

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

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.[31] 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,

[30] Whereas the logit model features a unique equilibrium (Milgrom and Roberts, 1990; Backus, 2014), the non-uniqueness of equilibria is a feature not only for the most flexible random coefficient model. It can also arise in other discrete choice settings, such as the logit model with conditional heteroskedasticity (Echenique and Komunjer, 2007).

[31] See Nevo (2000a) for a detailed explanation on 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. However, for mixed logit systems an equilibrium does not need to exist (Caplin and Nalebuff, 1991). This means that one should be careful when interpreting the counterfactual bounds, as it is under the assumption of equilibrium existence.[32] Also, if one cannot find a set of outcomes that satisfy the equilibrium bounds, conditional on the search being exhaustive, it necessarily implies that an equilibrium does *not* exist. Therefore, the method can provide numerical proof of the *non-existence* of an equilibrium at given parameter choices, but not the opposite.

5.1.1 Envelopes on Equilibrium Constraints

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

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

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

One can see that the 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.[33] 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, respectively. With this reformulation, a set of piece-wise linear constraints defines the demand-side equilibrium constraints.

[32] Aksoy-Pierson et al. (2013) provide conditions for equilibrium existence for a model with a discrete number of consumer types as the one considered here.

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

Additionally, to obtain aggregate shares, one can add the linear constraints,

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

To bound the supply-side equations, it is useful to use a similar technique 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 $\alpha_i = \alpha + \tilde{\alpha}_i$. One still needs to approximate these terms, as they are non-linear. Furthermore, they appear interacted with product markups $p_r - mc_r$ in the first order conditions.

To explain how to bound these terms, I focus on the own elasticity term. In such case, one needs to bound the 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 pieces. Because the shares were already used in the approximation of the log for the demand equations, this does not increase the computational costs excessively.

However, this function 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 that is being solved at each iteration.[34]

Finally, additional linear restrictions can be added that complement the above restrictions and can increase the precision of the algorithm. For example, one can also bound the constraints with:

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

Each of these additional constraints are linear, and can help construct tighter envelopes to the equilibrium conditions.

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

Table 4: Parameter Values for Example 3, Panel A

Demand Side		Supply Side	
Consumers types ns	5	Number of firms N	3
Coefficient β	$N([1.5, 1.5, .5],$ $Diag[.5, .5, .5])$	Characteristics X	$N([1, 1, 1],$ $[1, -.8, .3;$ $-.8, 1, .3;$ $.3, .3, 1;])$
Coefficient α	$N(3, .2)$	Marginal Cost	$[0, 0, 0]$
Coefficient μ	$N(0, .5)$	Coefficient ξ_j	$N(0, 1)$

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,

$$\max_{\mathbf{s}^*,\mathbf{p}^*} \text{ and } \min_{\mathbf{s}^*,\mathbf{p}^*} \quad \sum_i -w_i \gamma_{i0}$$

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

where I use $\mathbf{u}^*$ to define the additional integer variables that create the different regions of the piece-wise linear constraints.[35]

5.1.2 Example with Multiple Equilibria

Consider the parameter values in Table 4. One can approximate the solution to the demand system by computing piece-wise linear bounds to the logarithmic function between $(0, 1)$. Figure 6 shows an example of a piece-wise linear approximation, with ten pieces. Given the nature of the log function, it is useful to concentrate the pieces at low values of s.[36]

The results from the simulation are presented in Panel A of Table 5. The simulations consider five consumer types. One can observe that product 2 is particularly attractive under this parameterization, as it feature better characteristics. It also can charge larger markups. In terms of the approximation, one can observe that the upper and lower bounds to market shares are tight, suggesting that the log approximation induces limited loss in the precision of the bounds, as observed in columns four and five.

The table shows relatively tight bounds on shares and markups, suggesting the absence of multiple equilibria. Indeed, the example in Panel A exhibits a single equilibrium. Panel B presents an alternative example in which multiple equilibria occur. In this example, there are two products that are differentiated. Instead of having similar consumer types with varying sensitivities to price and product characteristics,

[35] See appendix for details, TBA.

[36] The figure has been computed with knots at $[10^{-5}, .01, .02, .04, .1, .2, .3, .4, .6, .8, 1]$. 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, I iterate over the relevant range of shares to improve the approximation.

Figure 6: Approximation of the function $\ln(s_{ij})$

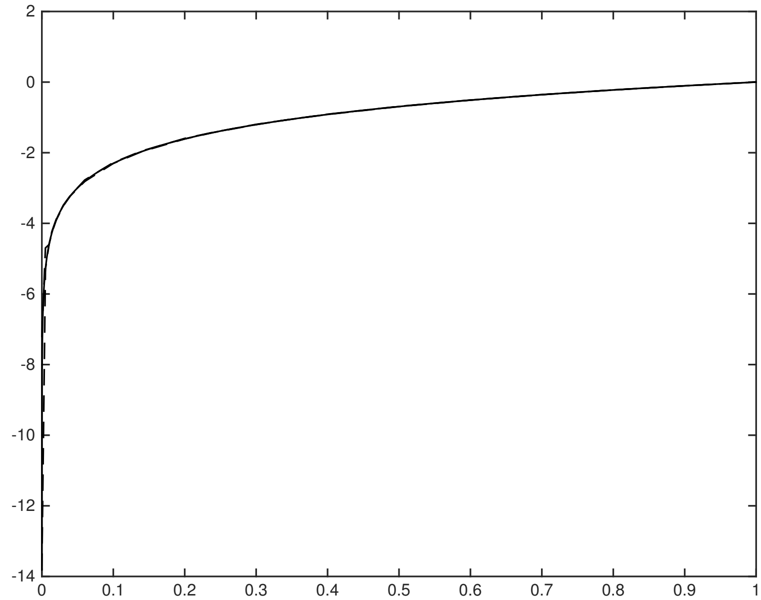

consumers also have substantial brand-specific taste. There exist a set of consumers that like product 1, a set of consumers that like product 2, and a third set of consumers that are more price-sensitive. This implies that a firm might face a trade-off: exploit the loyal consumers and set a high price, or lower the price and obtain demand from the price-sensitive consumers.

Panel B shows bounds on shares and markups for each of the two products under multiple equilibria. Whereas the iterated approach converges to a symmetric equilibrium in which both firms split market shares similarly and price at relatively low prices, the bounds approach characterizes more conservative bounds to market shares and markups.

It is important to note that Panel B in Table 5 presents bounds on market shares and markups that can be sustained in equilibrium for each of the products. However, not all bounds can be sustained at the same time (e.g., it is not possible for both Product 1 and 2 to enjoy a market share over 50%). The reason is that the methodology characterizes bounds for 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.

For this reason, it is useful to focus on bound the particular economic outcome one might be interested in, such as consumer surplus. As Panel B in Table 5 shows, the iterative methods converges to an equilibrium that is good for consumers. However, the bounds method highlights that other equilibria might arise in which consumers are worse off. This is already clear by noting that the bounds on the outside option range from 0.07 to 0.36, i.e., there are some equilibria in which quite a few consumers opt for the outside option.

5.2 Discrete Choice with Endogenous Attributes

The proposed methodology can also be extended to consider richer choice environments. For example, consider a setting in which firms choose both prices and product characteristics. Similar to the pricing game, multiple equilibria might arise. A regulator might be concerned that, after a merger, firms might

Table 5: Bounds on shares and markups (Example 1)

Panel A: Unique Equilibria

	Lower	Iter	Upper	Iter−Low	Upper−Iter
Shares					
Outside Option	0.22	0.22	0.22	3.3e-05	5.6e-05
Product 1	0.57	0.57	0.57	7.8e-05	6.2e-05
Product 2	0.21	0.21	0.21	5.6e-05	3.9e-05
Markups					
Product 1	1.18	1.18	1.18	1.3e-04	8.3e-05
Product 2	0.50	0.50	0.50	4.2e-05	4.2e-05
Consumer Surplus	0.00	1.89	0.00	-1.9e+00	1.9e+00

Panel B: Multiple Equilibria

	Lower	Iter	Upper	Iter−Low	Upper−Iter
Shares					
Outside Option	0.07	0.10	0.35	2.3e-02	2.6e-01
Product 1	0.31	0.45	0.51	1.5e-01	5.7e-02
Product 2	0.30	0.45	0.51	1.5e-01	6.4e-02
Markups					
Product 1	1.20	1.30	2.82	9.6e-02	1.5e+00
Product 2	1.21	1.32	3.06	1.1e-01	1.7e+00
Consumer Surplus	2.28	5.86	6.15	2.9e-01	3.6e+00

Notes: Results from Example 1. Panel A features a parameterization with two products and a unique equilibrium. Panel B features a parameterization with 2 products and 3 consumer types (loyal to product 1, loyal to product 2, and non-loyal), which features multiple equilibria.

decide to re-position their products in a way that is detrimental to consumers due to reduced competition.

Adding characteristic choices to the model can be achieved by introducing an additional vector of variables (characteristics of each car), together with a set of first-order conditions for characteristic choices. Product characteristics enter in consumer's utility as prices. The choice of characteristics has also 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. Similar to the pricing case, the logit formulation gives explicit solutions to $\partial s_r / \partial x_j$. These constraints can be trivially incorporated into the pricing and share constraints to compute equilibria with attribute choice.

The above examples exemplify how to compute bounds on BLP counterfactual outcomes (shares, markups, and attributes) in a robust fashion. Given the potential multiplicity of equilibria in these models, it can be a useful tool to be conservative about the impact of a given policy. This approach can be less appealing in settings in which there are many products or many consumer types, as the computational requirements can grow substantially. A scaled-downed version of a model can still be useful in situations in which iteration methods to computing equilibrium outcomes prove sensitive to starting values or fail to converge. It can also be used to provide ballpark starting values to 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 be useful to detect issues with the model parameterization (e.g., the lack of existence of an equilibrium in pure strategies).

6 Conclusions

I 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 constraint. This represents two innovations from the previous literature. First, I propose to use a relaxed constrained-optimization approach for equilibrium computation. Whereas constrained approaches have been typically used for estimation, or for computation of single agent games, I show how they can be used for computation purposes in strategic settings, as a way to 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 conservative 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, I have presented three examples in which counterfactual simulations might be difficult or multiplicity of equilibria might arise. First, I 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. Second, I 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, I presented a BLP model with multiple equilibria, and showed how counterfactual outcomes can be computed when endogenizing product choice.

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.

Backus, M. (2014). Uniqueness in N-Player Continuous Games with Application to Logit Pricing and Multi-Product Firms . Cornell University.

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

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.

Echenique, F. and Komunjer, I. (2009). Testing Models With Multiple Equilibria by Quantile Methods. *Econometrica*, 77(4):1281–1297.

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.

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.

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.

Ito, K. and Reguant, M. (2016). Sequential Markets, Market Power and Arbitrage. *American Economic Review*, forthcoming.

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.

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

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

Manski, C. F. (1989). Anatomy of the Selection Problem. *The Journal of Human Resources*, 24(3):343.

Manski, C. F. (1990). Nonparametric Bounds on Treatment Effects. *The American Economic Review*, 80(2):829–823.

Manski, C. F. (2003). *Partial Identification of Probability Distributions*. Springer Series in Statistics. Springer-Verlag, New York.

Manski, C. F. (2007). Partial identification of counterfactual choice probabilities. *International Economic Review*, 48(4):1393–1410.

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.

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.

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.

Tamer, E. (2010). Partial Identification in Econometrics. *Annual Review of Economics*, 2(1):167–195.

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

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.

Appendix

TBA – It will contain the mathematical programs characterizing the examples in the introduction (carbon regulation), example 1 (multi-unit auctions), example 2 (dynamic games), and other examples (e.g., BLP), with more implementation details. All examples are programmed in Python and use the solver Gurobi, which is available for free to the academic community.