Economics Letters 114 (2012) 190–194

Contents lists available at SciVerse ScienceDirect

Economics Letters

journal homepage: www.elsevier.com/locate/ecolet

A method for implementing counterfactual experiments in models with multiple equilibria

Victor Aguirregabiria *

Department of Economics, University of Toronto, 150 St. George Street, Toronto, ON, M5S 3G7, Canada

ARTICLE INFO

Article history:
Received 6 January 2011
Received in revised form
25 September 2011
Accepted 28 September 2011
Available online 6 October 2011

JEL classification:
C10
C60
L10

Keywords:
Structural models with multiple equilibria
Counterfactual experiments
Equilibrium selection
Homotopy methods

ABSTRACT

This paper proposes a homotopy method for implementing counterfactual experiments in empirical models with multiple equilibria. A key assumption is that the equilibrium selection function does not jump discontinuously between equilibria as we continuously change the structural parameters.

1. Introduction

Multiplicity of equilibria is a prevalent feature in static and dynamic games in economics. This indeterminacy poses practical problems when using these models for empirical analysis. Recent papers show that, under some assumptions, we can identify which equilibrium or equilibria is observed in the data, and then we can use two-step and sequential methods to estimate structural parameters in models with multiple equilibria.[1] Nevertheless, the indeterminacy problem associated with multiple equilibria still remains an issue when the researcher wants to use the estimated model to predict the effects of counterfactual changes in the structural parameters. Assumptions that identify the equilibrium played in the data are not enough to identify which equilibrium will be selected in a counterfactual scenario with values of the structural parameters that are different to the ones that we have estimated from the data.[2] Given that an attractive feature of structural models is the possibility of implementing counterfactual experiments, this is an important issue in structural econometrics.

This paper proposes a simple homotopy method for dealing with multiple equilibria when undertaking counterfactual experiments with an estimated model. The key assumption in our method is that the (unknown) equilibrium selection function does not jump discontinuously between equilibria as we change continuously the value of the structural parameters. In other words, the counterfactual equilibrium is of the same 'type' as the equilibrium in the data. The method proceeds in two steps. In the first step, and under the assumption that the equilibrium selection mechanism is a smooth function, we show how to obtain an approximation of the counterfactual equilibrium. The approximation in this first step may be inaccurate when the counterfactual experiment does not imply marginal changes in the parameters. Therefore, in the second step, and under the assumption that the counterfactual equilibrium is Lyapunov stable, we combine the first-step approximation with iterations in the equilibrium mapping. The idea is that the approximation obtained in the first step lies within the dominion of attraction of the counterfactual equilibrium.

* Tel.: +1 416 978 4358; fax: +1 416 978 6713.
E-mail address: victor.aguirregabiria@utoronto.ca.

1 See Aguirregabiria (2004), Aguirregabiria and Mira (2007), Bajari et al. (2007), Pakes et al. (2007), and Pesendorfer and Schmidt-Dengler (2008).

2 In some models, a possible approach to deal with this issue is to calculate all of the equilibria in the counterfactual scenario and then draw conclusions that are robust to whatever equilibrium is selected. Unfortunately, this approach is of very limited applicability in empirical applications where the different equilibria provide contradictory predictions of the effects we want to measure, as it is the case in dynamic games.

0165-1765/$ – see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.econlet.2011.09.041

We use a simple example to illustrate the differences between our approach and other simple methods, such as the selection of a counterfactual equilibrium that is closer (in Euclidean distance) to the equilibrium in the data, or equilibrium mapping iterations using the equilibrium in the data as the initial value. We compare the ability of these methods to satisfy the property that the same equilibrium 'type' is maintained in the factual and in the counterfactual scenarios.

2. Model

Let $\mathbf{x} \in X$ and $\mathbf{y} \in Y$ be two vectors of random variables with discrete and finite support.[3] Let $p_0(\mathbf{y}|\mathbf{x})$ be the *true* probability distribution function of $\mathbf{y}$ conditional to $\mathbf{x}$ in the population under study. We have a model for the conditional distribution p_0. The model is a parametric family of probability functions, $f(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta})$, where $\boldsymbol{\theta} \in \Theta$ is a vector of K parameters, and the set of parameters $\Theta \subset \mathbb{R}^K$ is compact. It is convenient to represent the probability functions $p_0(.|.)$ and $f(.|., \boldsymbol{\theta})$ as vectors in the Euclidean space. Define the vectors $\mathbf{P}_0 \equiv \{p_0(\mathbf{y}|\mathbf{x}) : (\mathbf{x}, \mathbf{y}) \in X \times Y\}$ and $F(\boldsymbol{\theta}) \equiv \{f(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta}) : (\mathbf{x}, \mathbf{y}) \in X \times Y\}$ that live in the set $[0, 1]^{|X||Y|}$ that we represent as $\mathcal{P}$. Let $\boldsymbol{\theta}_0$ be the true value of $\boldsymbol{\theta}$ in the population under study. We assume that the model is correctly specified and that $\boldsymbol{\theta}_0$ is identified from p_0, i.e., $\boldsymbol{\theta}_0$ is the unique value in the parameter space such that $p_0(\mathbf{y}|\mathbf{x}) = f(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta}_0)$ for every pair $(\mathbf{x}, \mathbf{y}) \in X \times Y$.

A key feature of the class of structural models that we consider in this paper is that the probability distribution $F(\boldsymbol{\theta})$ is not explicitly defined but it is only implicitly defined as *a* solution of an equilibrium or fixed point problem. Let $\Psi(\mathbf{P}, \boldsymbol{\theta})$ be a fixed-point or equilibrium mapping from $\mathcal{P} \times \Theta$ into $\mathcal{P}$ such that $\Psi(\mathbf{P}, \boldsymbol{\theta})$ is a vector-valued function $\{\psi(\mathbf{y}|\mathbf{x}, \mathbf{P}, \boldsymbol{\theta}) : (\mathbf{x}, \mathbf{y}) \in X \times Y\}$ and $\psi(.|., \mathbf{P}, \boldsymbol{\theta})$ is a conditional probability function that is twice continuously differentiable in $\mathbf{P}$ and $\boldsymbol{\theta}$. In models where the mapping $\Psi(., \boldsymbol{\theta})$ has a single fixed point for each value of $\boldsymbol{\theta}$, the vector or probability distribution $F(\boldsymbol{\theta})$ is unambiguously defined as *the* solution to the fixed point problem $\mathbf{P} = \Psi(\mathbf{P}, \boldsymbol{\theta})$. Instead, we consider here models with multiple equilibria. When the mapping $\Psi(., \boldsymbol{\theta})$ has multiple fixed points, the equilibrium restrictions do not uniquely characterize the probability distribution $F(\boldsymbol{\theta})$.

Suppose that the model has a finite number T of equilibrium 'types' that we index by $t \in \{1, 2, \ldots, T\}$. Let $\Pi_t(\boldsymbol{\theta})$ be the vector that represents equilibrium type t when the vector of parameters is $\boldsymbol{\theta}$, such that it satisfies the equilibrium restrictions $\Pi_t(\boldsymbol{\theta}) = \Psi(\Pi_t(\boldsymbol{\theta}), \boldsymbol{\theta})$. Note that some equilibrium types may exist only for a subset of points in the parameter space Θ. Under our conditions on the mapping $\Psi(\mathbf{P}, \boldsymbol{\theta})$, the equilibrium probability functions $\Pi_t(\boldsymbol{\theta})$ are continuous in $\boldsymbol{\theta}$. We assume that the probability distribution $F(\boldsymbol{\theta})$ is equal to $\Pi_{\tilde{t}(\boldsymbol{\theta})}(\boldsymbol{\theta})$ where $\tilde{t}(\boldsymbol{\theta}) \in \{1, 2, \ldots, T\}$ represents the *selected equilibrium* when the vector of parameters is $\boldsymbol{\theta}$. We can interpret function $\tilde{t}(\boldsymbol{\theta})$ as the equilibrium selection mechanism.

In the recent literature on empirical games, most papers do not make specific assumptions on the form of the equilibrium selection function $\tilde{t}(\boldsymbol{\theta})$. Here we also follow that approach. Then, we say that the model is *incomplete* in the sense that it does not provide a unique distribution of the endogenous variables for each possible value of the structural parameters.[4] We assume that our model is incomplete but identified, and that the researcher does not want to impose restrictions on the equilibrium selection mechanism in order to have a *complete* model. This class of econometric models includes as particular cases discrete models with social interactions (Brock and Durlauf, 2001), quantal response games (McKelvey and Palfrey, 1995), and static and dynamic games of incomplete information (Doraszelski and Satterthwaite, 2010), among others.

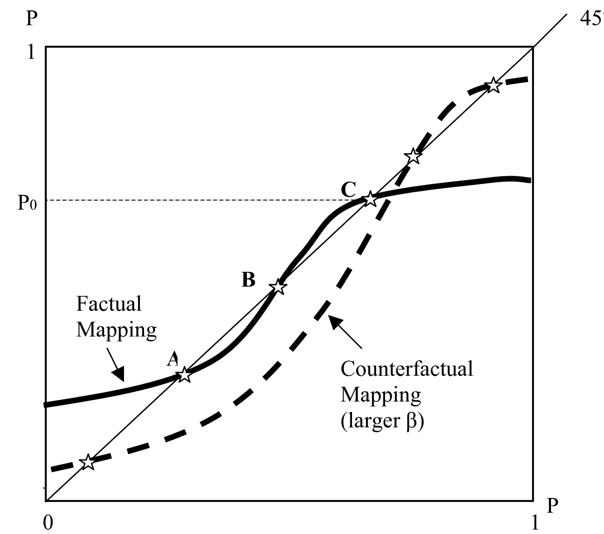

Fig. 1. Factual and counterfactual equilibrium mappings.

Suppose that $\mathbf{P}_0$ and $\boldsymbol{\theta}_0$ are point-identified given a random sample on $\mathbf{y}$ and $\mathbf{x}$. Let $\widehat{\boldsymbol{\theta}}_0$ and $\widehat{\mathbf{P}}_0$ be our consistent estimates of $\boldsymbol{\theta}_0$ and $\mathbf{P}_0$, respectively. The model establishes that $\widehat{\mathbf{P}}_0 = F(\widehat{\boldsymbol{\theta}}_0) = \Pi_{\tilde{t}(\widehat{\boldsymbol{\theta}}_0)}(\widehat{\boldsymbol{\theta}}_0)$, where the vector $\Pi_{\tilde{t}(\widehat{\boldsymbol{\theta}}_0)}(\widehat{\boldsymbol{\theta}}_0)$ is *a* fixed point of $\Psi(., \widehat{\boldsymbol{\theta}}_0)$. Remember that the vector $F(\widehat{\boldsymbol{\theta}}_0)$ is only implicitly defined as *a* solution to the fixed point problem. An important feature of the estimation methods that have been recently proposed and applied to estimate games with multiple equilibria is that the researcher does not have to compute the function(s) $\Pi_t(.)$. The researcher does not know the mapping $F(.)$, because he does not know neither the type-specific functions $\Pi_t(.)$ nor the equilibrium selection function $\tilde{t}(.)$. The researcher knows only a single point in the graph of the function $F(.)$, i.e., point $(\widehat{\boldsymbol{\theta}}_0, \widehat{\mathbf{P}}_0)$.

3. Counterfactual experiments

Let $\boldsymbol{\theta}_*$ be a value of the vector of structural parameters that is different to $\widehat{\boldsymbol{\theta}}_0$. We denote $\boldsymbol{\theta}_*$ as the vector of counterfactual values of the structural parameters. The researcher knows $\boldsymbol{\theta}_*$ and wants to obtain the counterfactual equilibrium $\mathbf{P}_*$ associated with $\boldsymbol{\theta}_*$, i.e., $\mathbf{P}_* = F(\boldsymbol{\theta}_*) = \Pi_{\tilde{t}(\boldsymbol{\theta}_*)}(\boldsymbol{\theta}_*)$. The researcher is interested in comparing this counterfactual equilibrium with the one estimated from the data. There are two main issues to obtain this counterfactual equilibrium. First, the researcher does not know the equilibrium functions $\Pi_t(\boldsymbol{\theta})$ for each equilibrium type t. And second, even if the researcher knew these functions, he does not know the equilibrium selection mechanism $\tilde{t}(\boldsymbol{\theta})$. He does not know which of the equilibria is selected when $\boldsymbol{\theta} = \boldsymbol{\theta}_*$. We need additional information/structure to select $\mathbf{P}_*$ from among the set of equilibria associated to $\boldsymbol{\theta}_*$.

We propose an approach that tries to impose minimum additional restrictions on the characteristics of the counterfactual equilibrium. Our approach is quite agnostic with respect to the equilibrium selection mechanism. We assume that there is such a mechanism, that it is a function, and that it does not "jump" between the possible equilibria when we move continuously over

[3] We describe our approach in the context of a class of models in which all the variables have a discrete and finite support. This is convenient because we can use standard derivatives to construct Taylor approximations. However, it is possible to extend this approach to models where variables have continuous support by using Banach spaces and Fréchet derivatives.

[4] There are good reasons why researchers may not want to impose additional restrictions on the equilibrium selection mechanism. First, incomplete models can be point-identified, and therefore assumptions on the equilibrium selection mechanism are not necessary for identification. And second, these additional restrictions, if they do not hold in the population under study, can induce biases in parameter estimates.

the parameter space. However, we do not specify any particular form for the equilibrium selection mechanism.

Assumption IET (Invariant Equilibrium Type): There is a convex subset $\mathcal{S}$ *of the parameter space* Θ *that includes* $\widehat{\boldsymbol{\theta}}_0$ *and* $\boldsymbol{\theta}_*$ *and where the equilibrium selection mechanism* $\tilde{t}(\boldsymbol{\theta})$ *is constant, i.e., for any* $\boldsymbol{\theta} \in \mathcal{S}$, $\tilde{t}(\boldsymbol{\theta}) = \tilde{t}(\widehat{\boldsymbol{\theta}}_0) = \tilde{t}(\boldsymbol{\theta}_*)$.

Assumption IET establishes that the equilibrium type does not change when we move from $\hat{\boldsymbol{\theta}}_0$ to $\boldsymbol{\theta}_*$. While this seems a strong assumption when evaluating a factual policy change, it is more reasonable when the researcher is interested in the prediction of the effects of a counterfactual new policy. When evaluating the effects of a factual policy change using data before and after the new policy, the data can identify the type of equilibrium that we have for each of the two subsamples. However, that is not the case when we want to predict the effects of a counterfactual change in parameters. In this situation, there is no data to identify the selected equilibrium after implementing the new policy. Therefore, we need to make some assumptions about equilibrium selection in the counterfactual scenario. It seems natural to start assuming that the policy change does not affect the selected equilibrium type.[5]

An implication of assumption IET is that the function $F(\boldsymbol{\theta}) = \Pi_{\tilde{t}(\boldsymbol{\theta})}(\boldsymbol{\theta})$ is continuously differentiable within the convex set $\mathcal{S}$ that includes $\widehat{\boldsymbol{\theta}}_0$ and $\boldsymbol{\theta}_*$.[6] Some equilibrium 'types' may disappear when we move along the parameter space. Assumption IET implicitly considers that the type of equilibrium $\tilde{t}(\widehat{\boldsymbol{\theta}}_0)$ does not disappear when we move from $\hat{\boldsymbol{\theta}}_0$ to $\boldsymbol{\theta}_*$. In principle, the researcher does not know if that is the case. However, we describe at the end of Section 4 a simple procedure to check for this departure from Assumption IET.

Example. Consider the following model within the class of models described above. P is a scalar. The equilibrium mapping Ψ is the function $G(\alpha + \beta h(P))$, where $G(.)$ is a continuously differentiable CDF of a random variable with support on the whole real line, α and β are parameters with $\beta > 0$, and $h(P)$ is a real-valued continuously differentiable function with $h'(P) > 0$. Let $\widehat{\boldsymbol{\theta}}_0 = (\widehat{\alpha}_0, \widehat{\beta}_0)$ and $\widehat{P}_0$ be consistent estimates of the true values of the parameters and of the factual equilibrium, respectively. Fig. 1 presents the estimated (or factual) equilibrium mapping $G(\widehat{\alpha}_0 + \widehat{\beta}_0 h(P))$, the estimated equilibrium $\widehat{P}_0$, and a counterfactual equilibrium mapping $G(\widehat{\alpha}_0 + \beta^* h(P))$. In that figure, an equilibrium is a value of P for which the curve meets with the 45° line. For the estimated equilibrium mapping there are three types of equilibria that we represent as A, B, and C. The values of these equilibria are $\Pi_A(\widehat{\boldsymbol{\theta}}_0)$, $\Pi_B(\widehat{\boldsymbol{\theta}}_0)$, and $\Pi_C(\widehat{\boldsymbol{\theta}}_0)$. Equilibria $\Pi_A(\widehat{\boldsymbol{\theta}}_0)$ and $\Pi_C(\widehat{\boldsymbol{\theta}}_0)$ are Lyapunov stable, i.e., the derivative $\partial G(\widehat{\alpha}_0 + \widehat{\beta}_0 h(P))/\partial P$ is smaller than 1 at $P = \Pi_A(\widehat{\boldsymbol{\theta}}_0)$ and at $P = \Pi_C(\widehat{\boldsymbol{\theta}}_0)$. Equilibrium $\Pi_B(\widehat{\boldsymbol{\theta}}_0)$ is not stable, i.e., the derivative $\partial G(\widehat{\alpha}_0 + \widehat{\beta}_0 h(P))/\partial P$ is greater than 1 at $P = \Pi_B(\widehat{\boldsymbol{\theta}}_0)$. Suppose that the data has been generated by equilibrium C such that $\widehat{P}_0 = F(\widehat{\boldsymbol{\theta}}_0) = \Pi_C(\widehat{\boldsymbol{\theta}}_0)$. The researcher is interested in the effect of a counterfactual change in the parameter β, keeping $\widehat{\alpha}_0$ constant. As shown in Fig. 1, a larger value of β implies an increase in the slope of the equilibrium mapping for every value of P. Also, in this example, an increase in β implies that the equilibrium function moves upwards for high values of P and downwards for low values of P. Let $\mathcal{S}$ be an interval in the real line such that: (i) the interval contains $\widehat{\beta}_0$; and (ii) type-C equilibrium exists for any β in this interval and keeping $\alpha = \widehat{\alpha}_0$. For the sake of illustration, we assume that equilibrium types A and B also exist for any β within $\mathcal{S}$. Fig. 2 presents the type-specific equilibrium functions $\Pi_A(\widehat{\alpha}_0, \beta)$, $\Pi_B(\widehat{\alpha}_0, \beta)$, and $\Pi_C(\widehat{\alpha}_0, \beta)$ over the interval $\mathcal{S}$. In this example, the probability function $F(\widehat{\alpha}_0, \beta)$ is equal to $\Pi_C(\widehat{\alpha}_0, \beta)$ for any value of β within $\mathcal{S}$.

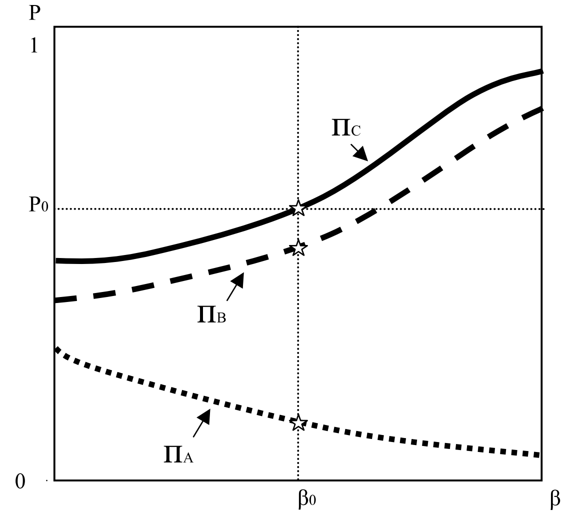

Fig. 2. Type-specific equilibria as a function of β.

4. A simple homotopy method

We want to obtain the counterfactual equilibrium associated with $\boldsymbol{\theta}_*$, that we denote $\mathbf{P}_*$. We know that $\widehat{\mathbf{P}}_0 = F(\widehat{\boldsymbol{\theta}}_0) = \Pi_{t_0}(\widehat{\boldsymbol{\theta}}_0)$ where $t_0 \equiv \tilde{t}(\widehat{\boldsymbol{\theta}}_0)$, and under Assumption IET, $\mathbf{P}_* = F(\boldsymbol{\theta}_*) = \Pi_{t_0}(\boldsymbol{\theta}_*)$. While $\widehat{\mathbf{P}}_0$, $\widehat{\boldsymbol{\theta}}_0$, and $\boldsymbol{\theta}_*$ are known to the researcher, $\mathbf{P}_*$ and the functions $F(.)$ and $\Pi_{t_0}(.)$ are unknown. The method proceeds in two steps.

Step 1 (*Taylor approximation*). Under Assumption IET, we can use a first order Taylor expansion to obtain an approximation to the counterfactual equilibrium $F(\boldsymbol{\theta}_*)$ around the estimated vector $\widehat{\boldsymbol{\theta}}_0$. We do not know the function $F(.)$ but it is possible to use the equilibrium condition to obtain the Jacobian matrix $\partial F(\widehat{\boldsymbol{\theta}}_0)/\partial \boldsymbol{\theta}'$ in terms of derivatives of the equilibrium mapping evaluated at $(\widehat{\mathbf{P}}_0, \widehat{\boldsymbol{\theta}}_0)$. A Taylor expansion of $F(\boldsymbol{\theta}_*)$ around $\widehat{\boldsymbol{\theta}}_0$ implies that:

$$P_* = F(\boldsymbol{\theta}_*) = F(\widehat{\boldsymbol{\theta}}_0) + \frac{\partial F(\widehat{\boldsymbol{\theta}}_0)}{\partial \boldsymbol{\theta}'}(\boldsymbol{\theta}_* - \widehat{\boldsymbol{\theta}}_0) + O(\|\boldsymbol{\theta}_* - \widehat{\boldsymbol{\theta}}_0\|^2). \quad (1)$$

Note that $F(\widehat{\boldsymbol{\theta}}_0) = \widehat{\mathbf{P}}_0$ that is known. Taking into account that $F(\widehat{\boldsymbol{\theta}}_0) = \Psi(F(\widehat{\boldsymbol{\theta}}_0), \widehat{\boldsymbol{\theta}}_0)$, differentiating this expression with respect to $\boldsymbol{\theta}$, and solving for $\partial F(\widehat{\boldsymbol{\theta}}_0)/\partial \boldsymbol{\theta}'$, we can represent this Jacobian matrix in terms of Jacobians of $\Psi(\mathbf{P}, \boldsymbol{\theta})$ evaluated at the estimated values $(\widehat{\mathbf{P}}_0, \widehat{\boldsymbol{\theta}}_0)$. That is,

$$\frac{\partial F(\widehat{\boldsymbol{\theta}}_0)}{\partial \boldsymbol{\theta}'} = \left(\mathbf{I} - \frac{\partial \Psi(\widehat{\mathbf{P}}_0, \widehat{\boldsymbol{\theta}}_0)}{\partial \mathbf{P}'}\right)^{-1} \frac{\partial \Psi(\widehat{\mathbf{P}}_0, \widehat{\boldsymbol{\theta}}_0)}{\partial \boldsymbol{\theta}'} \quad (2)$$

where $\mathbf{I}$ is the identity matrix, and Assumption IET implies that the matrix $\mathbf{I} - \partial \Psi(\widehat{\mathbf{P}}_0, \widehat{\boldsymbol{\theta}}_0)/\partial \mathbf{P}'$ is non-singular. Solving expression (2) into (1), we have that $F(\boldsymbol{\theta}_*) = \widetilde{\mathbf{P}}_* + O(\|\boldsymbol{\theta}_* - \widehat{\boldsymbol{\theta}}_0\|^2)$, where:

$$\widetilde{\mathbf{P}}_* \equiv \widehat{\mathbf{P}}_0 + \left(\mathbf{I} - \frac{\partial \Psi(\widehat{\mathbf{P}}_0, \widehat{\boldsymbol{\theta}}_0)}{\partial \mathbf{P}'}\right)^{-1} \frac{\partial \Psi(\widehat{\mathbf{P}}_0, \widehat{\boldsymbol{\theta}}_0)}{\partial \boldsymbol{\theta}'}(\boldsymbol{\theta}_* - \widehat{\boldsymbol{\theta}}_0). \quad (3)$$

Therefore, when $\|\boldsymbol{\theta}_* - \widehat{\boldsymbol{\theta}}_0\|^2$ is small, the vector $\widetilde{\mathbf{P}}_*$ provides a good approximation to the true counterfactual equilibrium $\mathbf{P}_*$. Note that

5 However, it is true that Assumption IET is more likely to hold for marginal policy changes (e.g., 10% increase in a license fee; or a 5% points increase in a sales tax) than for completely new policies. In fact, this caveat generally applies to most counterfactual experiments.

6 Following Doraszelski and Escobar (2010), it is possible to show that Assumption IET implies that for any vector of parameters $\boldsymbol{\theta}$ in the interior of the set $\mathcal{S}$, the equilibrium $F(\boldsymbol{\theta}) = \Pi_{\tilde{t}(\boldsymbol{\theta})}(\boldsymbol{\theta})$ is *regular*, and the Jacobian matrix $\mathbf{I} - \partial \Psi(F(\boldsymbol{\theta}), \boldsymbol{\theta})/\partial \mathbf{P}'$ is non-singular.

all the elements in the expression that describes $\widetilde{\mathbf{P}}_*$ are known to the researcher.

In some applications, the counterfactual experiments of interest are far from being marginal changes in the parameters. In such a situation, a first order Taylor approximation could be inaccurate. Higher-order approximations to $F(\boldsymbol{\theta}_*)$ can be used. It is possible to show that higher-order derivatives of $F(.)$ at $\widehat{\boldsymbol{\theta}}_0$ depend only on derivatives of Ψ at $(\widehat{\mathbf{P}}_0, \widehat{\boldsymbol{\theta}}_0)$, which are known to the researcher. However, in applications where the dimension of the vector $\mathbf{P}$ is large (e.g., in dynamic games with heterogeneous players), the numerical computation of high-order derivatives of Ψ with respect to $\mathbf{P}$ can be computationally very demanding. An alternative approach for improving the accuracy of our approximation is to combine the approach described above with iterations in the equilibrium mapping. That is Step 2 in our proposed method.

Step 2 (*Equilibrium mapping iterations*). Suppose that the equilibrium $\mathbf{P}_*$ is Lyapunov stable. This implies that there is a neighborhood of $\mathbf{P}_*$, say $\mathcal{N}$, such that if we iterate in the equilibrium mapping $\Psi(., \boldsymbol{\theta}_*)$ starting with a $\mathbf{P} \in \mathcal{N}$, then we converge to $\mathbf{P}_*$, i.e., if $\mathbf{P}_1 \in \mathcal{N}$ and $\mathbf{P}_{k+1} = \Psi(\mathbf{P}_k, \boldsymbol{\theta}_*)$ for any $k \geq 1$, then $\lim_{k\to\infty} \mathbf{P}_k = \mathbf{P}_*$ (Judd, 1998, Therorem 5.4.2). The neighborhood $\mathcal{N}$ is called the *dominion of attraction* of the stable equilibrium $\mathbf{P}_*$. Suppose that the Taylor approximation is precise enough such that $\widetilde{\mathbf{P}}_*$ belongs to the dominion of attraction of $\mathbf{P}_*$. Then, by iterating in the equilibrium mapping $\Psi(., \boldsymbol{\theta}_*)$ starting at $\widetilde{\mathbf{P}}_*$ we will obtain the counterfactual equilibrium $\mathbf{P}_*$.

5. An example

We use the example presented at the end of Section 3 to illustrate the method.[7] We also compare the proposed method with two alternative approaches for calculating a counterfactual equilibrium: (1) equilibrium mapping iteration using $\widehat{\mathbf{P}}_0$ as the initial value; and (2) equilibrium selection based on minimum Euclidean distance to $\widehat{\mathbf{P}}_0$. We compare these methods in terms of their ability to satisfy the property that the same equilibrium 'type' is maintained in the factual and in the counterfactual scenarios.

For the comparison of these methods in our example, it is important to take into account the dominion of attraction of each equilibrium type when iterating in the equilibrium mapping. Looking at Fig. 1, it should be clear that the dominion of attraction of an equilibrium type A is the interval $[0, \Pi_B)$. If we iterate in the equilibrium mapping starting at some probability P within the interval $[0, \Pi_B)$, we should converge to the stable equilibrium Π_A. Similarly, the dominion of attraction of an equilibrium type C is the interval $(\Pi_B, 1]$. And for non-stable equilibrium type B, the dominion of attraction is just the single point Π_B.

Consider Fig. 2 with the type-specific equilibria as functions of β over the interval $\mathcal{S}$. Given this figure, and taking into account the dominion of attraction of each equilibrium type, we can derive the predicted values for the counterfactual equilibrium P_* for every value β_* in the interval $\mathcal{S}$, and for each of the three methods. In Fig. 3 we present the prediction of our method. The schedule of predicted values $\widehat{P}_*$ is represented using the thick curve. The straight line tangent to curve $\Pi_C(\widehat{\alpha}_0, \beta)$ at point $(\widehat{\beta}_0, \widehat{P}_0)$ represents the first order Taylor approximation around $(\widehat{\beta}_0, \widehat{P}_0)$, that we denote $\widetilde{P}_*$.[8] When this straight line is above the

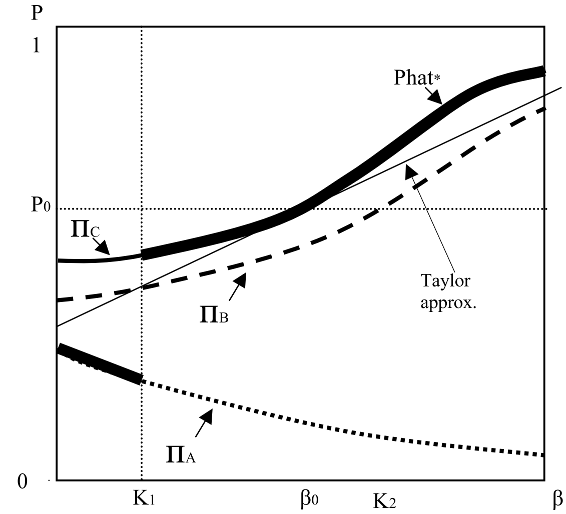

Fig. 3. Predicted counterfactual equilibrium proposed method (as a function of β).

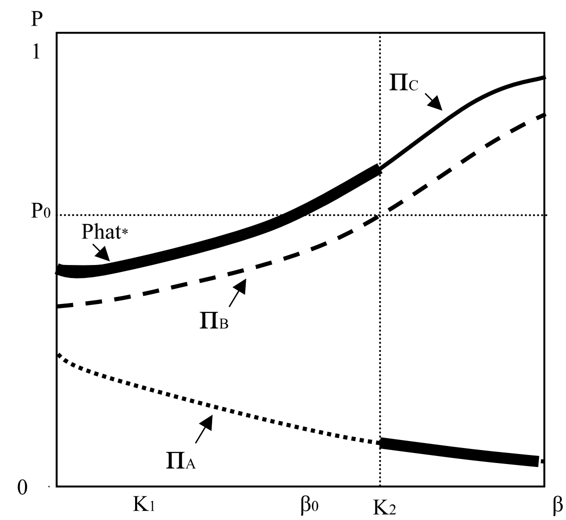

Fig. 4. Predicted counterfactual equilibrium (as a function of β). Method: Iterations starting from factual equilibrium.

curve for equilibrium type Π_B (i.e., $\widetilde{P}_* > \Pi_B$) we have that the Taylor approximation lies in the dominion of attraction of equilibrium type C, and therefore Step-2 of the method provides a counterfactual $\widehat{P}_* = \Pi_C(\widehat{\alpha}_0, \beta_*)$ that satisfies the smoothness condition. The straight line for the Taylor approximation crosses curve Π_B at point K_1. For values of β smaller than K_1 we have that $\widetilde{P}_* < \Pi_B$, such that the Taylor approximation lies in the dominion of attraction of equilibrium type A, and the prediction of this method is $\widehat{P}_* = \Pi_A(\widehat{\alpha}_0, \beta_*)$ that does not satisfy the smoothness condition. For values of β_* within the interval $\mathcal{S}$ and smaller than K_1, the Taylor approximation (first step) provides a better approximation to the counterfactual equilibrium than the combination of Steps 1 and 2.

Fig. 4 presents the predicted counterfactual equilibrium when using equilibrium mapping iterations initialized at $\widehat{P}_0$. For values of β_* such that $\widehat{P}_0 > \Pi_B(\widehat{\alpha}_0, \beta_*)$, we have that $\widehat{P}_0$ lies in the dominion of attraction of equilibrium type C, and therefore the prediction is $\widehat{P}_* = \Pi_C(\widehat{\alpha}_0, \beta_*)$ that satisfies the smoothness condition. When $\widehat{P}_0 < \Pi_B(\widehat{\alpha}_0, \beta_*)$, we have that $\widehat{P}_0$ lies in the dominion of attraction of equilibrium type A and the prediction of this method is $\widehat{P}_* = \Pi_A(\widehat{\alpha}_0, \beta_*)$ that does not satisfy the smoothness condition. It is important to note that this method provides different predictions than our method. Each method may

7 Aguirregabiria and Ho (in press) have applied this method to implement counterfactual experiments in an empirical dynamic game of oligopoly competition of the airline industry.

8 In this example, this first order approximation has the following form:

$$\widetilde{P}_* = \widehat{P}_0 + \frac{g(\widehat{\alpha}_0 + \widehat{\beta}_0 h(\widehat{P}_0))}{1 - \widehat{\beta}_0 h(\widehat{P}_0) g(\widehat{\alpha}_0 + \widehat{\beta}_0 h(\widehat{P}_0))} (\beta_* - \widehat{\beta}_0)$$

where $g(.)$ is the density function of the CDF $G(.)$.

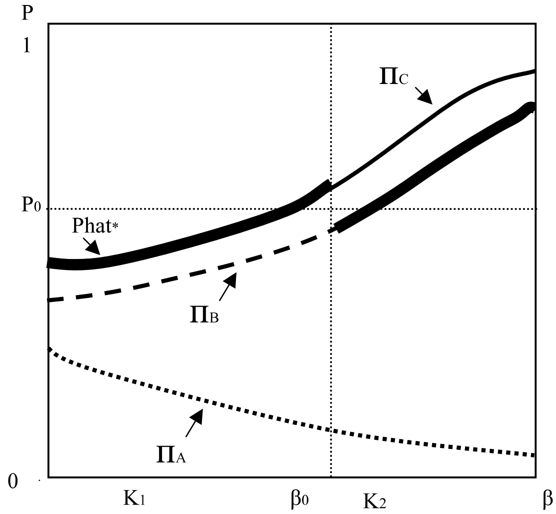

Fig. 5. Predicted counterfactual equilibrium (as a function of β). Method: Minimum distance to factual equilibrium.

not satisfy the smoothness condition for different range of values of the counterfactual parameters.

Finally, Fig. 5 presents the predicted counterfactual equilibrium when using minimum distance to $\widehat{P}_0$ as the equilibrium selection criterion.[9] We see that it also fails to satisfy the smoothness condition, and it does it for a different range of values than the other two methods.

This example also illustrates that there is a trade-off when using 'Step-2' (i.e., equilibrium mapping iterations) in our method. A method that applies only the Taylor approximation in 'Step-1' is more conservative in the sense that it always satisfies the smoothness condition at the cost of providing, "typically", a worse approximation. In contrast, the application of 'Step-2' "typically" provides a better approximation but it involves the potential risk of switching to a different equilibrium type and not satisfying the smoothness condition.

Fig. 3 (and 4) suggests a straightforward extension of our proposed method (though with additional computational cost) that can deal with the potential problem of switching to a different equilibrium type when applying Step 2. Instead of calculating the counterfactual equilibrium only for the counterfactual value β_*, we can calculate the schedule of counterfactual equilibria for a fine grid of points within the interval $[\widehat{\beta}_0, \beta_*]$. Then, we can check whether this schedule is a continuous function of β.

Acknowledgments

I have benefited from conversations with Ron Borkovski, Arvind Magesan, and Pedro Mira. I would like to thank the Social Science and Humanities Research Council of Canada (SSHRCC) for financial support.

References

Aguirregabiria, V., 2004. Pseudo maximum likelihood estimation of structural models involving fixed-point problems. Economics Letters 84, 335–340.
Aguirregabiria, V., Ho, C.-Y., 2011. A dynamic oligopoly game of the US airline industry: estimation and policy experiments. Journal of Econometrics (in press). http://dx.doi.org/10.1016/j.jeconom.2011.09.013.
Aguirregabiria, V., Mira, P., 2007. Sequential estimation of dynamic discrete games. Econometrica 75, 1–53.
Bajari, P., Benkard, L., Levin, J., 2007. Estimating dynamic models of imperfect competition. Econometrica 75, 1331–1370.
Brock, W., Durlauf, S., 2001. Discrete choice with social interactions. The Review of Economic Studies 68, 235–260.
Doraszelski, U., Escobar, J., 2010. A theory of regular Markov perfect equilibria in dynamic stochastic games genericity, stability, and purification. Theoretical Economics 5, 369–402.
Doraszelski, U., Satterthwaite, M., 2010. Computable Markov-perfect industry dynamics. Rand Journal of Economics 41, 215–243.
Judd, K., 1998. Numerical Methods in Economics. The MIT Press, Cambridge, Massachusetts.
McKelvey, R., Palfrey, T., 1995. Quantal response equilibria for normal form games. Games and Economic Behavior 10, 6–38.
Pakes, A., Ostrovsky, M., Berry, S., 2007. Simple estimators for the parameters of discrete dynamic games (with entry/exit examples). Rand Journal of Economics 38, 373–399.
Pesendorfer, M., Schmidt-Dengler, P., 2008. Asymptotic least squares estimators for dynamic games. The Review of Economic Studies 75, 901–928.

9 Note that this method requires one to calculate all the equilibria of the mapping $\Psi(., \boldsymbol{\theta}_*)$ and then selecting the equilibrium with the smallest Euclidean distance to $\widehat{\mathbf{P}}_0$. In general, this method will be computationally more costly to implement that the other two methods.