Economics Bulletin

Volume 37, Issue 4

Selecting Equilibria using Best-Response Dynamics

Vincent Boucher
Universite Laval

Abstract

I propose a simple simulation procedure for large games with multiple equilibria. The simulation procedure is based on a best-response dynamic. The implied equilibrium selection mechanism is intuitive: more stable equilibria are selected with higher probability.

I would also like to thank Charles Bellemare, Luc Bissonnette, Arnaud Dufays, Michel Roland, as well as the other participants to the internal seminar series at Université Laval, for interesting discussions and comments. I also gratefully acknowledge financial support from the Fonds de Recherche du Québec - Société et culture and from the Social Sciences and Humanities Research Council.
Citation: Vincent Boucher, (2017) "Selecting Equilibria using Best-Response Dynamics", *Economics Bulletin*, Volume 37, Issue 4, pages 2728-2734
Contact: Vincent Boucher - vincent.boucher@ecn.ulaval.ca.
Submitted: June 19, 2017. **Published:** December 01, 2017.

1 Introduction

Equilibrium refinements are often used when games feature many Nash equilibria. The choice of the appropriate refinement is motivated by its theoretical properties and its ability to predict observed behavior (e.g. Govindan et al. (2005), Kim (1996)). However, in large games, most equilibrium refinements are not feasible due to a curse of dimensionality (e.g. in network formation models, see Chandrasekhar (2015)).

In this paper, I present a simple, simulation-based, equilibrium selection mechanism (i.e. equilibrium refinement) based on a best-response dynamic. I focus on generalized ordinal potential (GOP) games (Monderer & Shapley, 1996) with finite strategy spaces.

I provide that a simple and intuitive simulation procedure, based on a best-response dynamic. I show that the likelihood of an observation under this simulation procedure is equal to its expected relative stability (which is formally defined in section 2.1).

I present two simple examples. The first one is a two-by-two game and allows to compare this new equilibrium selection mechanism with alternative mechanisms. The second one is a network formation game and allows to see how the curse of dimensionality is alleviated.

In the next section, I present the game, the main results, as well as the two examples. I conclude in section 3 by discussing potential extensions.

2 The Game

Consider a finite set N of n individuals, each of whom chooses a strategy $a_i \in A_i$ and has preferences represented by the utility u_i. I assume that A_i is finite for all i, and denote $A = \times_{i\in N} A_i$ and $a = (a_1, ..., a_n) \in A$. I consider the following random utility model:

$$u_i(a; \varepsilon) \tag{1}$$

where ε has a known probability density function f. Note that in most cases, $u_i(a; \varepsilon)$ may have detailed structure such as $u_i(a; \varepsilon) = u(a_i, a_{-i}; x_i, \theta, \varepsilon)$, where x_i is some series of observed characteristics and θ is a parameter of interest. I denote the set of Nash equilibria (NE) of the game $\Gamma_\varepsilon = \langle A_i, u_i(\cdot; \varepsilon)\rangle_{i=1}^n$ by $A^*_\varepsilon \subseteq A$.

For simplicity, I assume no indifference, in the sense that for all $i \in N$, whenever $a_i \neq a'_i$, we have that $u_i(a; \varepsilon) \neq u_i(a'_i, a_{-i}; \varepsilon)$. Note that this assumption can be replaced by an arbitrary tie-breaking rule. Note also that in many interesting cases, this will hold almost surely since f is absolutely continuous. An example is provided in section 2.2

I'm interested in the empirical content of Γ_ε, which is summarized by its likelihood function. When the game has more than one equilibrium (for a positive measure of ε), the likelihood must be completed by an *equilibrium selection mechanism* (Tamer, 2003; Galichon & Henry, 2011). An equilibrium selection mechanism is a probability distribution π with support on A^*_ε. The likelihood is therefore:

$$\mathcal{L}(a) = \int \pi(a|A^*_\varepsilon)\mathbb{1}\{a \in A^*_\varepsilon\} f(\varepsilon) d\varepsilon$$

In essence, this is the likelihood that a is a NE of Γ_ε, *and* that it is the selected equilibrium among all equilibria in A^*_ε.[1]

In this paper, I present an intuitive equilibrium selection mechanism, based on a simple simulation procedure. The procedure goes as follows, for a fixed ε.

Algorithm 1 (Uniform Best-Response Dynamic)

1. *Draw uniformly* $a \in A$. *Draw uniformly a permutation* λ *on* N. *Set* $t = 0$.
2. *Sequentially for all* i *and according to the order in* λ, *let:*
    - $a_i^{t+1} = \arg\max_{a_i} u_i(a_i, a_{-i}^t; x, \varepsilon)$ *and*
    - $a_{-i}^{t+1} = a_{-i}^t$.

[1] Note that the indicator function can be omitted since π has support on A^*_ε.

Let $t \to t+1$.

3. *Repeat step 2 until $a^t = a^{t+n}$.*

As shown in section 2.3, even if the cardinality of A is large, uniform sampling may not be computationally intensive. Also, provided that this simulation procedure is well defined and meaningful (see section 2.1), one can easily perform well-known simulation-based inference methods, such as (in a classical setting) simulated General Method of Moments, or simulated Maximum Likelihood (see Gourieroux & Monfort (1996)), or (in a Bayesian setting) approximate Bayesian computation (see Marin et al. (2012)).

I now describe the properties of the uniform best-response dynamic.

2.1 Results

I now study the properties of the likelihood function implied by the uniform best-response dynamics. The first result follows directly from literature, and ensures that the procedure is well defined.

I assume that the game has a *generalized ordinal potential* (Monderer & Shapley, 1996), i.e. a function $Q(a;\varepsilon)$, such that, for all $i \in N$ and any $a_i, a_i' \in A_i$,

$$u_i(a_i, a_{-i};\varepsilon) - u_i(a_i', a_{-i};\varepsilon) > 0 \text{ implies } Q(a_i, a_{-i};\varepsilon) - Q(a_i', a_{-i};\varepsilon) > 0.$$

For generalised ordinal potential games, best-response dynamics never get stuck in infinite loops.

Proposition 1 (Monderer & Shapley (1996)) *If the game has a generalized ordinal potential, then the simulation procedure converges to a Nash equilibrium (NE) in finite time.*

Of course, if the procedure converges, it necessarily converges to a NE, by construction. One of the properties of GOP games is that the procedure always converges. This convergence happens in finite time since the strategy space is finite. Monderer & Shapley (1996) also show that if all *better-response paths* converge (i.e. not just the ones based on best-response for a specific order of play), then the game is necessarily a GOP. This therefore implies that there is little hope for the convergence of the uniform best-response dynamic for non-GOP games.

I now turn to the attractiveness of the uniform best-response dynamic as an equilibrium refinement. I need to introduce some additional definitions. Given ε and λ, an *improvement path* from a to a', is a finite sequence of strategy profiles $a^1, ..., a^m \in N$ such that $a^1 = a$ and $a^m = a'$, and such that for (a^k, a^{k+1}), there is a unique individual i_k such that $a^{k+1}_{-i_k} = a^k_{-i_k}$ and $a^{k+1}_{i_k} = \arg\max_{a_{i_k}} u_i(a_{i_k}, a^k_{-i_k};\varepsilon)$, where i_k follows the sequence of play λ. That is $i_k = \lambda(k)$ and $i_{n+k} = \lambda(k)$ for $k \leq n$. I will let $z^*(a, a';\varepsilon,\lambda) = 1$ if a' can be reached from a using some improvement path, given λ and ε. I let $z^*(a, a';\varepsilon,\lambda) = 0$ otherwise. One can easily see that any improvement path stops once it reaches a NE.

For each NE $a^*_\varepsilon \in A^*_\varepsilon$, the *basin of attraction* is defined as the set of strategy profiles which necessarily lead to a^*_ε.

$$B^*(a^*_\varepsilon;\varepsilon,\lambda) = \{a' \in A : z^*(a', a^*_\varepsilon;\varepsilon,\lambda) = 1\}$$

Note that any NE belongs to its own basin of attraction, i.e. $B(a^*_\varepsilon;\varepsilon,\lambda) \ni a^*_\varepsilon$. Note also that since it is impossible to have an improvement path from one NE to another, $B(a^*_\varepsilon;\varepsilon,\lambda)$ only contains one NE.

This notion of basin of attraction is quite standard, for instance, in the literature on evolutionary dynamics (see, Ellison (2000) among others). Here, for convenience, I will extend the notion of basin of attraction to all strategy profiles by letting $B^*(a;\varepsilon,\lambda) = \emptyset$ for all $a \in A \setminus A^*_\varepsilon$, so the basin of attraction of a is non-empty iff $a \in A^*_\varepsilon$.

I define the *stability index* of a strategy profile as follows

Definition 1 *Fix ε, then for any strategy profile $a \in A$, its stability index (SI) is given by:*

$$SI^*(a;\varepsilon) = \sum_\lambda \#B^*(a;\varepsilon,\lambda).$$

In essence, $SI(a)$, for a strategy profile a gives the number of combinations of strategy profiles and sequences of play such that the best-response dynamics ends at a. We have the following:

Figure 1: A Two-by-Two Example

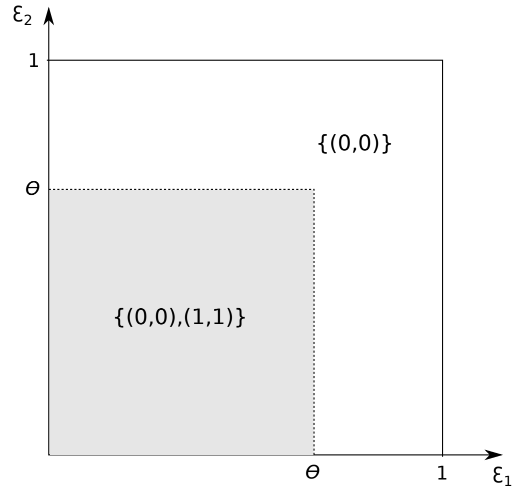

Proposition 2 (Main Result) *Under the uniform best-response dynamic, the likelihood of each strategy profile is proportional to its expected stability index:*

$$\mathcal{L}(a) = \frac{1}{C}\int SI^*(a;\varepsilon)f(\varepsilon)d\varepsilon$$

where $C = \sum_{a'} SI^*(a;\varepsilon) > 0$ *is independent of* ε*.*

See Appendix for a proof.

An interesting special case is when there exists a unique stable equilibrium $a^s_\varepsilon \in A^*_\varepsilon$ such that $B^*(a^*_\varepsilon;\varepsilon,\lambda) = A\setminus(A^*_\varepsilon \setminus \{a^s_\varepsilon\})$, for all λ. That is, a^*_ε attracts all improving paths that did not originate from another NE. In this special case, the unique stable equilibrium is selected with probability going to 1 as $|A| \to \infty$ (provided that the set of equilibria is bounded).

I now present two simple examples.

2.2 A Two-by-Two Example

Consider the following (somewhat classical) game.[2] There are two individuals with strategies $a_i \in \{0,1\}$, $i = 1,2$. Their utility are as follows for $i,j = 1,2$: $u_i(a) = (\theta a_j - \varepsilon_j)a_i$, where ε_i and ε_j are distributed according to the cumulative distribution F, with full support on $[0,1]^2$, and where $\theta \in (0,1]$ is an exogenous parameter. The following matrix summarizes the game:

	0	**1**
0	$0,0$	$0,-\varepsilon_1$
1	$-\varepsilon_2,0$	$\theta-\varepsilon_2,\theta-\varepsilon_1$

This game has (potentially) two NE: $a^* = (0,0)$, irrespective of the value of ε, and $a^* = (1,1)$ if $\varepsilon \in [0,\theta]^2$ (see Figure 1).

This game has many interesting properties. It is supermodular (Topkis, 1979), a (non-symmetric) coordination game, as well as a potential game (Monderer & Shapley, 1996). A potential game is a special case of GOP, where $Q(a;\varepsilon)$ is such that

$$u_i(a_i,a_{-i};\varepsilon) - u_i(a'_i,a_{-i};\varepsilon) = Q(a_i,a_{-i};\varepsilon) - Q(a'_i,a_{-i};\varepsilon).$$

[2]The game was first proposed by Jovanovic (1989), and more recently used by Tamer (2003) and Galichon & Henry (2011), among others.

Table 1: Alternative Equilibrium Selection Mechanism

Equilibrium Selection Mechanism	$\mathcal{L}(0,0)$	$\mathcal{L}(1,1)$
Risk-dominance	$1-\hat{F}(\theta)$	$\hat{F}(\theta)$
Maximum of $Q(a)$	$1-\hat{F}(\theta)$	$\hat{F}(\theta)$
Payoff-dominance	$1-F(\theta,\theta)$	$F(\theta,\theta)$

One can easily check that Q is also a generalized ordinal potential. The potential function for this game is:

$$Q(a_1,a_2)=\theta a_1 a_2-\varepsilon_1 a_2-\varepsilon_2 a_1.$$

As shown by Ui (2001) and Carbonell-Nicolau & McLean (2014), the maximum of the potential function offers an attractive equilibrium refinement. Indeed, this simple game allows for many natural equilibrium selection mechanism. Let $\hat{F}$ be the cumulative distribution of $\varepsilon_1+\varepsilon_2$. Table 1 shows the likelihood of the two NE, under alternative equilibrium refinements.

I now look at the likelihood implied by the uniform best-response dynamic. The basins of attraction (for the two potential NE) are as follows:

- $B^*((0,0);[0,\theta]^2,(1,2))=\{(0,0),(1,0)\}$
- $B^*((0,0);[0,\theta]^2,(2,1))=\{(0,0),(0,1)\}$
- $B^*((0,0);[0,1]^2\setminus[0,\theta]^2,\lambda)=\{(0,0),(0,1),(1,0),(1,1)\}$ for any λ
- $B^*((1,1);[0,\theta]^2,(1,2))=\{(0,1),(1,1)\}$
- $B^*((1,1);[0,\theta]^2,(2,1))=\{(1,0),(1,1)\}$

Then, the stability indices (for the two potential NE) can be computed:

- $SI^*((0,0);[0,\theta]^2)=4$
- $SI^*((0,0);[0,1]^2\setminus[0,\theta]^2)=8$
- $SI^*((1,1);[0,\theta]^2)=4$

while the sum is: $\sum_a SI^*(a;\varepsilon)=8$, for all values ε.

The likelihood function is therefore:

$$\begin{aligned}\mathcal{L}(0,0)&=\frac{4}{8}F(\theta,\theta)+\frac{8}{8}(1-F(\theta,\theta))=1-F(\theta,\theta)/2\\ \mathcal{L}(1,1)&=\frac{4}{8}F(\theta,\theta)+\frac{0}{8}(1-F(\theta,\theta))=F(\theta,\theta)/2\end{aligned}$$

Note that in the special case where $(\varepsilon_1,\varepsilon_2)$ are independently and uniformly distributed on $[0,1]$, we have that $F(\theta,\theta)=\theta^2$ and $\hat{F}(\theta)=\frac{\theta^2}{2}$. The likelihood for the uniform best-response dynamic is therefore the same as the likelihood of the maximum of the potential function, and risk-dominance.

I now present an example of a game typically affected by the curse of dimensionality.

2.3 Network Formation

There is a population of $n\geq 3$ individuals, each of whom has a strategy $a_i\in A_i=\{0,1\}^{n-1}$. The strategy profile $A=\times_i A_i$ represents a (directed) network and has a cardinality of $2^{n(n-1)}$. See Chandrasekhar (2015) for a discussion of the empirical challenges of network formation games.

In order to implement the uniform Best-response dynamic, one has to sample uniformly from a set of $2^{n(n-1)}$ networks (i.e. from the set of strategy profiles A). This can be done in two steps.

1. Draw k from a binomial distribution $B(k;n(n-1),1/2)$.

2. Create each link independently with probability $p = k/n(n-1)$.

The second step is simply an Erdos-Reni random network (Erdös & Rényi, 1959), which draws uniformly among networks with $pn(n-1)$ links (as n becomes large). To understand why the first step draws the number of links k coherently, note the number of networks with k links is given by the binomial coefficient: $\begin{pmatrix} n(n-1) \\ k \end{pmatrix}$. That is, the number of possibilities for having k links in a population of $n(n-1)$ pairs. However, as n increases, drawing from a binomial distribution is computationally intensive. Fortunately, a simple (and very precise) approximation is to draw from a normal distribution with mean $n(n-1) - 1/2$ and variance $n(n-1)/4$ (e.g. Schader & Schmid (1989)).

Therefore, combining steps 1 and 2 allows to sample uniformly from A. I now conclude with some remarks.

3 Remarks and Extensions

This paper presents a simple, and intuitive equilibrium selection mechanism. It, however, relies on two important assumptions, namely that the sampling procedure is uniform, and that the strategy space is finite.

The simulation procedure *can* be applied to more general strategy spaces. For example, to compact strategy space. In such a case, even if the simulation procedure may not converge in finite time, it converges to an ϵ-NE in finite time. (Monderer & Shapley, 1996) Therefore, in practice, the simulation procedure could be used to simulate relatively good approximations of the games's NE (given some tolerance threshold).

Also, here, the sampling of the initial strategy profile, and of the order of play is assumed to be uniform. If this is intuitive (and necessary for the result of proposition 2 to hold), the simulation procedure can easily be extended to any distribution, potentially depending of a parameter of interest. In section 2.2 for instance, one could assume that one of the two individuals (assuming it is identifiable) plays first with probability $\beta \in (0, 1)$. Similarly, in section 2.3, the initial network could be selected using a random network formation model (e.g. Pin & Rogers (2016)).

Overall, the method is simple, flexible and intuitive and could be applied to a wide range of games and economic applications.

4 References

Carbonell-Nicolau, O., & McLean, R. P. (2014). Refinements of nash equilibrium in potential games. *Theoretical Economics*, *9*(3), 555–582.

Chandrasekhar, A. (2015). Econometrics of network formation. In *The Oxford Handbook of the Economics of Networks*, (pp. 303–357).

Ellison, G. (2000). Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution. *The Review of Economic Studies*, *67*(1), 17–45.

Erdös, P., & Rényi, A. (1959). On random graphs, i. *Publicationes Mathematicae (Debrecen)*, *6*, 290–297.

Galichon, A., & Henry, M. (2011). Set identification in models with multiple equilibria. *The Review of Economic Studies*, (p. rdr008).

Gourieroux, C., & Monfort, A. (1996). *Simulation-based econometric methods*. Oxford university press.

Govindan, S., Wilson, R., et al. (2005). Refinements of nash equilibrium. In *The New Palgrave Dictionary of Economics, 2nd Edition*. Citeseer.

Jovanovic, B. (1989). Observable implications of models with multiple equilibria. *Econometrica: Journal of the Econometric Society*, (pp. 1431–1437).

Kim, Y. (1996). Equilibrium selection inn-person coordination games. *Games and Economic Behavior*, *15*(2), 203–227.

Marin, J.-M., Pudlo, P., Robert, C. P., & Ryder, R. J. (2012). Approximate bayesian computational methods. *Statistics and Computing*, *22*(6), 1167–1180.

Monderer, D., & Shapley, L. S. (1996). Potential games. *Games and economic behavior*, *14*(1), 124–143.

Pin, P., & Rogers, B. (2016). Stochastic network formation and homophily. In Y. Bramoullé, A. Galeotti, & B. W. Rogers (Eds.) *The Oxford Handbook on the Economics of Networks*, chap. 7. Oxford: Oxford University Press.

Schader, M., & Schmid, F. (1989). Two rules of thumb for the approximation of the binomial distribution by the normal distribution. *The American Statistician*, *43*(1), 23–24.

Tamer, E. (2003). Incomplete simultaneous discrete response model with multiple equilibria. *Review of Economic Studies*, *70*(1), 147–165.

Topkis, D. M. (1979). Equilibrium points in nonzero-sum n-person submodular games. *SIAM Journal on Control and Optimization*, *17*(6), 773–787.

Ui, T. (2001). Robust equilibria of potential games. *Econometrica*, *69*(5), 1373–1380.

5 Appendix: Proof of Proposition 2

Fix ε. It is sufficient to show that

$$\pi(a|A^*_\varepsilon)\mathbb{1}\{a \in A^*_\varepsilon\} = \frac{SB_\varepsilon(a)}{\sum_{a'} SB_\varepsilon(a)}.$$

If a is not a NE , then both expressions are equal (to 0). If a is a NE, then we have:

$$\pi(a|A^*_\varepsilon) = \frac{SB_\varepsilon(a)}{\sum_{a'} SB_\varepsilon(a)}.$$

Fix λ, then I show that it is possible to partition the set of strategy profiles using the basins of attraction. Under proposition 1, it is sufficient to show that if $a \in B^*(a';\varepsilon,\lambda)$, then $a \notin B^*(a'';\varepsilon,\lambda)$. Suppose otherwise. Since the improvement path must follow the order of play λ, it implies that, starting from the same strategy profile, and using the same sequence of play, the improvement paths would lead to different NE. This is impossible since preferences are strict.

We therefore have: $B^*(a;\varepsilon,\lambda) \cap B^*(a';\varepsilon,\lambda) = \emptyset$ for all $a \neq a'$. This implies that $\pi(a|A^*_\varepsilon,\lambda) = \#B^*(a;\varepsilon,\lambda)/\#A$ since starting strategy profiles are drawn uniformly.

Then, we have:

$$\pi(a|A^*_\varepsilon) = \frac{1}{n!}\sum_\lambda \pi(a|A^*_\varepsilon,\lambda) = \frac{1}{n!}\frac{1}{\#A}\sum_\lambda \#B^*(a;\varepsilon,\lambda)$$

or equivalently:

$$\pi(a|A^*_\varepsilon) = \frac{1}{n!}\frac{1}{\#A}SI^*(a;\varepsilon)$$

And since $\pi(a|A^*_\varepsilon)$ is a probability, we must have: $\sum_{a'} SI^*(a';\varepsilon) = \#A \cdot n!$, which completes the proof.[3]
QED

[3] Note that alternatively, one can see that $\sum_{a'} SI^*(a';\varepsilon) = \sum_\lambda \sum_a \#B^*(a;\varepsilon,\lambda) = \sum_\lambda \#A = n!\#A$.