


June 2010

Evolutionary Stability of First Price Auctions

Fernando Louge and Frank Riedel



EVOLUTIONARY STABILITY OF FIRST PRICE AUCTIONS

Fernando Louge Frank Riedel

Institute of Mathematical Economics
Bielefeld University

June 29, 2010

Abstract

This paper studies the evolutionary stability of the unique Nash equilibrium of a first price sealed bid auction. It is shown that the Nash equilibrium is not asymptotically stable under payoff monotonic dynamics for arbitrary initial populations. In contrast, when the initial population includes a continuum of strategies around the equilibrium, the replicator dynamic does converge to the Nash equilibrium. Simulations are presented for the replicator and Brown–von Neumann–Nash dynamics. They suggest that the convergence for the replicator dynamic is slow compared to the Brown–von Neumann–Nash dynamics.

JEL Classification: C73, D44

1 Introduction

Auctions have been extensively studied using classical game theory equilibrium concepts.¹ On the other hand, surprisingly little is known about the out-of-equilibrium properties of these auctions.² Since there are many predictions of the standard theory that are not observed in practice,³ out-of-equilibrium dynamics might provide some valuable insight about real world auctions.

This note presents a new approach to first price auctions. It consists in applying tools from the theory of evolution in games with continuous strategies. The main objective is to analyze the stability properties of the Nash equilibrium under the replicator dynamic as well as general payoff-monotonic dynamics.⁴

We will consider a standard model of evolution, the continuous time replicator dynamic and its generalizations, payoff-monotonic dynamics. The literal interpretation of

¹See for example Krishna (2002) and Milgrom (2004) for surveys.

²Exceptions include some learning models as the ones of Hon-Snir, Monderer, and Sela (1998) and Saran and Serrano (2007).

³See Kagel and Levin (2008) for a survey.

⁴The cornerstone solution concept in evolutionary game theory is the Evolutionarily Stable Strategy (ESS). However, in a game with infinitely many strategies, as it is the case with auctions, it is by now well known that ESS is not sufficient to guarantee dynamic stability (see for example Oechssler and Riedel (2002), Cressman, Hofbauer, and Riedel (2006)). Analogous concepts have been introduced for continuous strategy games, like the Continuously Stable Strategy (CSS) of Eshel (1983) and the Neighborhood Invader Strategy (NIS) of Apaloo (1997).

evolutionary reproduction and fitness does not apply to auctions, of course; however, many learning algorithms that are more plausible for human interactions lead to such dynamics as well, in particular those based on imitation⁵. We thus believe that our results provide valuable insights into the robustness of the unique Nash equilibrium of first-price auctions for boundedly rational agents.

Our main result is a negative one: the unique Nash equilibrium fails to be asymptotically stable under regular payoff-monotone dynamics. This might come as a surprise as the Nash equilibrium is not only unique but also strict. However, we find strategies that are able to invade the Nash equilibrium. Even though the invader strategies are not best reply to themselves (i.e. they are not equilibria), they fare better against themselves than the equilibrium strategy.

We prove our result with the help of a static stability concept, Neighborhood Invader Stability, or NIS. We show that our Nash equilibrium fails to be NIS. Similarly to a result of Eshel and Sansone (2003) and Cressman (2005), we then show that the Bayesian equilibrium is NOT asymptotically stable under replicator dynamics (and similar regular payoff-monotone dynamics). This is a negative result which shows that, maybe, first price auctions are more difficult to play for boundedly rational players. Indeed, we cannot expect convergence to the equilibrium from arbitrary initial conditions.

Our negative result begs the question: how do evolutionary and learning dynamics in first-price auctions look like? We shall show next that for “many” initial distributions, in particular those with full support, the bidding behavior is close to the Nash prediction. As long as evolutionary deviations are sufficiently dispersed, we can thus expect to see bidding behavior close to the Bayesian Nash prediction, even though convergence to the exact equilibrium cannot be expected.

One reason for this might be that – when restricting the strategy space to a suitably parametrized class – the equilibrium does satisfy another static stability concept, namely continuous stability, or CSS. In a first price auction, this condition is good enough to ensure convergence to the Nash equilibrium for many (but not all!) initial distributions close to the Nash equilibrium.⁶

Finally, we present some simulations based on different initial conditions. There are two main reasons for running them. First, they allow to see the behavior of some dynamics for which there are very few theoretical results, such as the BNN dynamic. Second, theoretical results are asymptotic in nature, but there is very little known about the evolutionary process in the medium run.

In these simulations we find that the replicator dynamic converges to the equilibrium if the initial distribution is sufficiently well dispersed as predicted by the theory, but convergence takes place at a very slow rate. The Brown-von Neumann-Nash dynamic converges much faster in our simulations.

The paper is organized as follows. Section 2 presents the model. The theoretical results regarding convergence are found in Section 3. Section 4 include the simulations of the model done for the replicator and BNN dynamics. Concluding remarks are offered in Section 5.

⁵See Schlag (1998). Sandholm (2010) includes several different imitative models that lead to the replicator dynamic.

⁶Louge (2010) presents some conditions that guarantee the convergence to CSS.

2 The model

We consider a single object for sale through a first price sealed bid auction. There are 2 bidders, 1 and 2. The valuation for bidder i is v_i . Valuations are drawn independently from the uniform distribution on the interval $[0, 1]$. Player i bids according to measurable function $s_i : [0, 1] \rightarrow \mathbb{R}_+$. The player with the highest bid wins the auction and the pays his own bid. In case of a tie, each bidder wins the object with equal probability.

It is well known that the unique symmetric Nash equilibrium of this game is $s_1(v) = s_2(v) = s^*(v) = \frac{v}{2}$.⁷ This equilibrium is strict, so it implies that s^* is an ESS.

We consider a restricted game where players choose strategies of the form

$$s(\theta, v) = \frac{v^\theta}{2} \quad (1)$$

for some parameter $\theta > 0$. In this restricted game, the strategy set of players is $\Theta = \mathbb{R}_{++}$ and the unique Nash equilibrium is to play $\theta^* = 1$ for both players. Let $S(\Theta)$ denote the set of bidding functions in the restricted game. Notice that $S(\Theta)$ allows for both overbidding (when $\theta < 1$) and underbidding (when $\theta > 1$).

We compute next the ex ante payoff.

Lemma 1. *The payoff function for bidder i against bidder j in the restricted game is*

$$\pi(\theta_i, \theta_j) = \frac{\theta_i \theta_j (1 + 2\theta_j)}{2\theta_i^2 + 2\theta_i^2 \theta_j + 6\theta_i \theta_j + 4\theta_i \theta_j^2 + 4\theta_j^2} \quad (2)$$

Proof. See the Appendix. □

Lemma 1 shows that even a simple class of bidding function can result in a payoff function more complex than the standard quadratic payoffs used in the evolution literature.

We will consider a standard model of evolution, the continuous time replicator dynamic. This dynamic is probably the most studied in the literature, introduced originally by Taylor and Jonker (1978) for biological reproduction in finite games. It is very popular in economics as well, since it is the outcome of appealing models of imitative behavior.⁸

At time $t \in [0, \infty)$, the state is given by a population $P_t \in \Delta\Theta$. Let $\Pi(P, Q)$ denote the expected payoff of population P against population Q . The excess payoff is defined as

$$\sigma(\theta, P) = \Pi(\delta_\theta, P) - \Pi(P, P) \quad (3)$$

where δ_θ is the Dirac distribution on θ .

For an initial population P_0 , the trajectory P_t^R evolves according to the replicator dynamic if

$$\frac{dP^R}{dt}(A) = \int_A \sigma(\theta, P^R) P^R(d\theta) \quad (4)$$

where A is a set in the Borel σ -algebra of Θ . Oechssler and Riedel (2001) have shown that (4) is well defined and has a unique solution.

⁷In general, the equilibrium strategy in a first price auction is given by $b_i = E[\max_{j \neq i} v_j | v_i]$. See Krishna (2002).

⁸See Schlag (1998). Sandholm (2010) includes several different imitative models that lead to the replicator dynamic.

The replicator dynamic can be generalized to the class of payoff monotonic dynamics. We will use the definition introduced in Oechssler and Riedel (2002). Let $\hat{\Pi}$ be the average conditional payoff on a Borel set A with positive probability $P(A) > 0$:

$$\hat{\Pi}(A, P) = \frac{1}{P(A)} \int_A \Pi(\delta_\theta, P) P(d\theta) \quad (5)$$

A dynamic is called *payoff monotonic* if for all strategy sets A, B with $P(A), P(B) > 0$,

$$\hat{\Pi}(A, P) > \hat{\Pi}(B, P) \text{ if and only if } \frac{1}{P(A)} \frac{dP}{dt}(A) > \frac{1}{P(B)} \frac{dP}{dt}(B) \quad (6)$$

3 Dynamic (In)Stability

In this section we will state our main results regarding the stability of the Nash equilibrium of the auction. We will consider initial populations that contain θ^* in their support. As it was argued by Oechssler and Riedel (2002), the relevant notion of distance is the Prohorov metric.

3.1 Instability

When considering all possible initial populations, we obtain a negative result: the Nash equilibrium is not asymptotically stable under any payoff monotonic dynamic.

Theorem 1. *The unique Bayesian Nash equilibrium in first-price auctions with uniform types is not asymptotically stable under payoff monotonic dynamics, in particular the replicator dynamic.*

Proof. We prove this result by showing first that the Nash equilibrium fails to be NIS in our restricted game. Note that for instability, it is enough to find one possible invader. Our parametrization is thus not a loss of generality.

Recall that θ^* is NIS if there exists an $\varepsilon > 0$ such that for all θ with $|\theta - \theta^*| < \varepsilon$

$$\pi(\theta^*, \theta) > \pi(\theta, \theta) \quad (7)$$

We can thus quickly check this condition.

$$\pi(1, \theta) - \pi(\theta, \theta) = \frac{\theta}{2(1+2\theta)} - \frac{1+2\theta}{6(2+\theta)} \quad (8)$$

$$= -\frac{(1-\theta)^2}{6(2+\theta)(1+2\theta)} < 0 \quad (9)$$

The condition (7) for a NIS fails for all $\theta \neq \theta^*$.

Fix $\varepsilon > 0$. Consider any initial distribution P with support on $\{\theta^*, \theta\}$, for some θ such that $|\theta - \theta^*| < \varepsilon$. Then,

$$\rho(\delta_{\theta^*}, P) = \min\{p^*, |\theta - \theta^*|\} < \varepsilon \quad (10)$$

where ρ is the Prohorov distance and $p^* = P(\{\theta^*\})$. The average conditional payoffs are

$$\hat{\Pi}(\{\theta^*\}, P) = p^* \pi(\theta^*, \theta^*) + (1 - p^*) \pi(\theta^*, \theta) \quad (11)$$

$$\hat{\Pi}(\{\theta\}, P) = p^* \pi(\theta, \theta^*) + (1 - p^*) \pi(\theta, \theta) \quad (12)$$

From (9), there exists $\bar{p} \in (0, 1)$ such that if $p^* < \bar{p}$, then $\hat{\Pi}(\{\theta^*\}, P) < \hat{\Pi}(\{\theta\}, P)$. Since the dynamic is payoff monotonic, we obtain

$$\frac{1}{p^*} \frac{dp^*}{dt} < \frac{1}{1-p^*} \frac{d(1-p^*)}{dt} \quad (13)$$

Since $p^* \in (0, 1)$, it follows that $\frac{d}{dt}p^* < 0$. Therefore, δ_{θ^*} is not asymptotically stable under the replicator dynamic. $\square$

The proof exploits the fact the equilibrium fails to satisfy the NIS criterion. If a strategy is not NIS, then it is not asymptotically stable under replicator dynamics, as it is shown in Eshel and Sansone (2003) and Cressman (2005). Theorem 1 provides a general argument for payoff monotonic dynamics. Since the equilibrium is not a NIS, then there exists an invader that is close enough to the equilibrium such that the equilibrium is not locally superior. If the initial population is dimorphic and the weight on the invading strategy is high enough, then the equilibrium obtains a lower conditional average payoff. Under the class of payoff monotonic dynamics, the population drifts away from the equilibrium.

3.2 Stability in the restricted game

The argument presented in the previous section relies of an invading population with finite support. We find that a positive result can be obtained by imposing some conditions on the initial population. In particular, if the support of the population is an interval (that includes the equilibrium strategy θ^*) then the replicator dynamic converges to the equilibrium. One reason for this is that the Bayesian Nash equilibrium does satisfy a certain static stability condition, CSS.

Recall that $\theta^* \in \Theta$ is called continuously stable, or CSS, if it is evolutionarily stable (ESS) and there exists an $\varepsilon > 0$ such that for all θ with $|\theta - \theta^*| < \varepsilon$ there exists $\eta > 0$ so that for all $\theta' \in \Theta$ with $|\theta - \theta'| < \eta$

$$\pi(\theta', \theta) > \pi(\theta, \theta) \text{ if and only if } |\theta - \theta^*| > |\theta' - \theta^*|. \quad (14)$$

Proposition 1. *In the restricted game, the Bayesian Nash equilibrium $\theta^* = 1$ is CSS.*

Proof. See the Appendix. $\square$

The next step is to notice that our restricted game exhibits strategic complementarities. It is straightforward that the best response function BR is increasing, where

$$BR(\theta) = \frac{\sqrt{2}\theta}{\sqrt{1+\theta}}. \quad (15)$$

We can obtain then the following result.

Proposition 2. *In the restricted game, if the support of the population is an interval that contains the equilibrium strategy θ^* , then the replicator dynamic converges to the Bayesian Nash equilibrium in the Prohorov metric.*

The result Proposition 2 follows from Loue (2010), where an argument based on iterated dominance shows that CSS are stable in games with strategic complements if the support of the population is a continuum.

The only qualification to this positive result is that considers only the game restricted to strategy set $S(\Theta)$. One could also consider the game restricted to a different class of bidding functions. We have also considered linear bidding functions, obtaining the same results⁹.

4 Simulations for Replicator and BNN Dynamics

In this section we present some simulations using different initial distributions. These were conducted for the replicator dynamic as well as the BNN dynamic. We start with the definition of the BNN dynamic. The positive part of the excess payoffs is

$$\sigma_+(\theta, P) = \max\{\sigma(\theta, P), 0\} \quad (16)$$

A trajectory P_t^B evolves according to the BNN dynamic if

$$\frac{dP^B}{dt}(A) = \int_A \sigma_+(\theta, P^B) d\theta - P^B(A) \int_{\Theta} \sigma_+(\theta, P^B) d\theta \quad (17)$$

for all A in the Borel σ -algebra of Θ .

We ran simulations of the replicator and the BNN dynamics for the set $S(\Theta)$ of bidding functions. The strategy space considered is $\Theta = [\frac{1}{2}, \frac{3}{2}]$ and it was discretized in 101 equidistant points.

The differential equations were approximated by the Euler method, with a time step d . In other words, for dynamic $N = R, B$ we have,

$$P_{t+d}^N(\theta_i) = P_t^N(\theta_i) + d \times P_t^N(\theta_i) \times \frac{dP^N}{dt}(\theta_i) \quad (18)$$

The simulation was iterated 1,000,000 times with time step $d = 0.1$. Below we show the distributions for each dynamic after 100,000 and 1,000,000 iterations.

Figure 1 is obtained from a uniform initial distribution. We can see that both dynamics appear to converge to the Nash equilibrium, although the replicator dynamic appears to be much slower. Interestingly, the mean strategy coincides with the Nash equilibrium at the initial distribution, but immediately becomes larger than one for both dynamics (recall that $\theta > 1$ corresponds to underbidding). This jump is much larger for the replicator dynamic.

Figure 2 from an initial distribution that puts .99 of probability on $\theta = 1.05$ and divides the remaining probability equality among all other strategies. In this case, the behavior of our two dynamics is dramatically different. The BNN dynamic moves the distribution gradually towards the equilibrium. It can also be seen that the mean strategy also moves towards $\theta^* = 1$, although not as fast as expected. For the replicator dynamic, the distribution appears not to move. We know, however, that the initial distribution is not a restpoint. In the fourth graph in can be seen that the probability of $\theta = 1.04$ is increasing, but remaining always a very small magnitude. This shows that convergence predicted by the theory is taking place, but at a very small rate.

Figure 3 considers a normal distribution with mean 1.05 and standard deviation 0.1. Similarly to Figure 1, it appears that both dynamics move towards the equilibrium, with the BNN dynamics being the fastest.

⁹The results are available from the authors upon request.

Figure 4 considers an initial distribution that puts .99 of probability (divided equally) on two strategies: $\theta_1 = 0.9$ and $\theta_2 = 1.1$. The remaining probability is distributed uniformly for all other strategies. In this case, replicator and BNN behave differently. The BNN dynamic evolves toward a unimodal distribution centered around the Nash equilibrium fairly fast. The mean strategy is never far from $\theta^* = 1$. On the other hand, the replicator dynamic shifts very quickly all the weight from $\theta_1 = 0.9$ towards $\theta_2 = 1.1$. Once this happens, the distribution starts a slow convergence path towards the equilibrium.

Figure 5 considers an initial distribution that puts .99 of probability (divided equally) on three strategies: $\theta_1 = 0.9$, $\theta^* = 1$ and $\theta_2 = 1.1$. For this initial distribution, both dynamics appear to converge. The replicator dynamic, however, puts relatively more weight in $\theta_2 = 1.1$ for the first hundred thousand periods, although this asymmetry is corrected after that.

5 Conclusion

This note introduces a new approach to the study of evolution in auctions models. The most general result obtained is a negative one: for an arbitrary initial population, the unique Bayes Nash equilibrium of a first price is unstable under payoff monotonic dynamics. Since this is an instability result, the restriction of the strategy space is without loss of generality.

In contrast, our second result is that when the initial population satisfy some regularity condition (namely, full support around the equilibrium), then the replicator dynamic converges to the Nash equilibrium of the auction. This result is qualified by the fact that we consider only a class of bidding functions. Moreover, the same result holds with different restrictions (e.g. linear bidding functions).

We consider a restriction of the strategy space in order to apply results from the theory of evolution in continuous strategy games. While ideally the set of bidding functions considered should be unrestricted, this note provides a first step in that direction.

Finally, some simulations are presented for different initial populations. The following conclusions can be drawn. First, it confirms our convergence result for the replicator dynamic, even when the initial population does not include a continuous interval around the equilibrium but a grid. Second, it shows that the BNN appears to have even better convergence properties. This is valuable since there are very few theoretical results for the BNN dynamic. Third, it highlights the issue of the speed of convergence. It somewhat qualifies our positive result for the replicator dynamic, since the time needed for convergence is surprisingly large in some cases.

This note suggests several paths for future research. First, the approach introduced here can be used in a large number of economic applications, starting from other auctions formats in particular to private information models in general where strategies are functions from an infinite set of types. Second, more progress needs to be done by considering a less restrictive strategy space.

References

APALOO, J. (1997): "Revisiting strategic models of evolution: The concept of neighborhood invader strategies," *Theoretical Population Biology*, 52, 7177.

- CRESSMAN, R. (2005): “Stability of the Replicator Equation with Continuous Strategy Space,” *Mathematical Social Sciences*, 50, 127–147.
- (2009): “Continuously Stable Strategies, Neighborhood Superiority and Two-Player Games with Continuous Strategy Space,” *International Journal of Game Theory*, 38, 221–247.
- CRESSMAN, R., J. HOFBAUER, AND F. RIEDEL (2006): “Stability of the replicator equation for a single-species with a multi-dimensional continuous trait space,” *Journal of Theoretical Biology*, 239, 273–288.
- ESHEL, I. (1983): “Evolutionary and continuous stability,” *Journal of Theoretical Biology*, 103, 99111.
- ESHEL, I., AND E. SANSONE (2003): “Evolutionary and dynamic stability in continuous population games,” *Journal of Mathematical Biology*, 46, 445459.
- HON-SNIR, S., D. MONDERER, AND A. SELA (1998): “A Learning Approach to Auctions,” *Journal of Economic Theory*, 82, 65–88.
- KAGEL, J., AND D. LEVIN (2008): “Auctions: A Survey of Experimental Research, 1995–2008,” forthcoming *Handbook of Experimental Economics*, Vol. 2.
- KRISHNA, V. (2002): *Auction Theory*. Academic Press.
- LOUGE, F. (2010): “On The Stability Of CSS Under The Replicator Dynamic,” mimeo, Bielefeld University.
- MILGROM, P. R. (2004): *Putting Auction Theory to Work*. Cambridge University Press.
- OECHSSLER, J., AND F. RIEDEL (2001): “Evolutionary dynamics on infinite strategy spaces,” *Economic Theory*, 17, 141–162.
- (2002): “On the Dynamic Foundation of Evolutionary Stability in Continuous Models,” *Journal of Economic Theory*, 107, 223–252.
- SANDHOLM, W. H. (2010): *Population Games and Evolutionary Dynamics*. MIT Press, forthcoming.
- SARAN, R., AND R. SERRANO (2007): “The Evolution of Bidding Behavior in Private-Values Auctions and Double Auctions,” mimeo.
- SCHLAG, K. (1998): “Why Imitate, and if so, How? A Boundedly Rational Approach to Multi-Armed Bandits,” *Journal of Economic Theory*, 78, 130–156.
- TAYLOR, P., AND L. JONKER (1978): “Evolutionary stable strategies and game dynamics,” *Mathematical Biosciences*, 40, 145–156.

Appendix and Figures

Proof of Lemma 1. Since bidding functions are strictly increasing in v , the probability of a tie is zero. Therefore, the ex-ante payoffs for a bidder using strategy s_i against strategy s_j are given by

$$\pi(s_i, s_j) = \int_0^1 [v_i - s_i(v_i)] \Pr[s_i(v_i) > s_j(v_j) \mid v_i] dv_i \quad (19)$$

According to (1), bidder i wins the auction against j whenever $v_j^{\theta_j} < v_i^{\theta_i}$. It follows then

$$\pi(\theta_i, \theta_j) = \int_0^1 \left[v - \frac{v^{\theta_i}}{2} \right] v^{\frac{\theta_i}{\theta_j}} dv \quad (20)$$

$$= \frac{1}{\frac{\theta_i}{\theta_j} + 2} - \frac{1}{2 \left(\frac{\theta_i}{\theta_j} + \theta_i + 1 \right)} \quad (21)$$

$$= \frac{\theta_i \theta_j (1 + 2\theta_j)}{2\theta_i^2 + 2\theta_i^2 \theta_j + 6\theta_i \theta_j + 4\theta_i \theta_j^2 + 4\theta_j^2} \quad (22)$$

□

Proof of Proposition 1. We will use a result from Cressman (2009). Namely, an ESS θ^* is a CSS if for all $\varepsilon > 0$ it is risk dominant in the 2×2 game with strategies $\{\theta^*, \theta\}$, where $|\theta - \theta^*| < \varepsilon$.

In other words, s^* is a CSS if for all θ close to the equilibrium,

$$RD(\theta) = [\pi(\theta^*, \theta^*) - \pi(\theta, \theta^*)] - [\pi(\theta, \theta) - \pi(\theta^*, \theta)] > 0 \quad (23)$$

From Lemma 1, we obtain

$$RD(\theta) \equiv \left[\frac{1}{6} - \frac{3\theta}{2(2+\theta)(1+2\theta)} \right] - \frac{(1-\theta)^2}{6(2+\theta)(1+2\theta)} \quad (24)$$

$$= \frac{(1-\theta)^2}{6(2+\theta)(1+2\theta)} > 0 \quad (25)$$

Therefore, θ^* is risk dominant against any strategy $\theta \neq \theta^*$. □




Figure 1: Simulation for a uniform initial distribution




Figure 2: Simulation for an initial distribution that puts .99 of probability on $\theta = 1.05$ and its uniform in the rest of the strategy space




Figure 3: Simulation for a normal initial distribution with mean 1.05 and standard deviation 0.1




Figure 4: Simulation for an initial distribution that puts .99 of probability (divided equally) on two strategies: $\theta_1 = 0.9$ and $\theta_2 = 1.1$




Figure 5: Simulation for an initial distribution that puts .99 of probability (divided equally) on three strategies: $\theta_1 = 0.9$, $\theta^* = 1$ and $\theta_2 = 1.1$