

Complementarities between Mathematical and Computational Modeling: The Case of the Repeated Prisoners' Dilemma

*Xavier Vilà**

Universitat Autònoma de Barcelona

Abstract

We study the properties of the well known Replicator Dynamics when applied to a finitely repeated version of the Prisoners' Dilemma game. We characterize the behavior of such dynamics under strongly simplifying assumptions (i.e. only 3 strategies are available) and show that the basin of attraction of defection shrinks as the number of repetitions increases. After discussing the difficulties involved in trying to relax the “strongly simplifying assumptions” above, we approach the same model by means of simulations based on genetic algorithms. The resulting simulations describe a behavior of the system very close to the one predicted by the replicator dynamics without imposing any of the assumptions of the mathematical model. Our main conclusion is that mathematical and computational models are good complements for research in social sciences. Indeed, while computational models are extremely useful to extend the scope of the analysis to complex scenarios hard to analyze mathematically, formal models can be useful to verify and to explain the outcomes of computational models.

Keywords: Computational Economics, Model-To-Model Analysis, Genetic Algorithms, Evolutionary Game Theory, Prisoners' Dilemma

JEL-Codes: C63, C73

1 Introduction

In the growing field of Agent-Based computer simulations applied to social sciences, model replication is considered a key issue. Indeed, asserting whether the observed results of a particular simulation of a model are correct or generalizable is a difficult task when no formal (i.e. mathematical) proof is provided. Only replication, comparison, alignment, and other related techniques can shed some light on the validity of simulations. See Axelrod (1997) for a methodological motivation on this and Hegselmann & Will (2008), Will (2009), and Macy & Sato (2010) for an inspirational debate. The

*Financial support from grant ECO2008-04756 (Grupo Consolidado-C) from the Spanish Ministry of Science and Innovation and from grant SGR2009-578 of the Generalitat de Catalunya are gratefully acknowledged. Departament d'Economia i d'Història Econòmica and MOVE. Universitat Autònoma de Barcelona. Xavier.Vila@uab.eu

work by Izquierdo *et al.* (2009), for instance, shows that the alignment of computational and mathematical models can “assist researchers in understanding the dynamics of simulation models” (Izquierdo *et al.*, 2009). Our work contains one such comparison. We put side-by-side two different analysis (mathematical and computational) of the same model: the evolution of strategies in the repeated prisoners’ dilemma.

We first consider, based on Imhof *et al.* (2005), the case in which the evolutionary system is described by a deterministic dynamic system that uses expected values. Using strong simplifying assumptions the model can be solved, and a complete description of how the process behaves is provided.

The second approach, based on Miller (1996), is a computational simulation in which finite automata are used to represent the strategies played, and a decentralized adaptive process based on the models of genetic algorithms simulates the stochastic evolutionary process. With this technique we can relax some of the strong assumptions used in the first approach and still obtain the same basic results.

We like to think that the limitations of the first approach (mathematical) provide a good motivation for the second approach (Computer-Based simulations). Indeed, although both approaches address the same problem, we show that the use of computational techniques allows us to relax hypothesis and overcome the limitations of the mathematical approach. On the other hand, it is shown that the mathematical model is extremely useful in order to explain the behavior and the causality of the results of the computational model

The choice of the repeated prisoners’ dilemma to conduct the experiment described above is not arbitrary. It is a well know and largely studied game, and many things about it have been learned thanks to the tools of formal game theory. But when the game is studied from an evolutionary perspective, the results are not always clear. The findings in Boyd & Lorberbaum (1987) and Binmore & Samuelson (1992), for instance, show that evolutionary stable solutions may fail to exists in many versions of the game. On the other hand, experiments and simulations like those reported in Axelrod (1985), Axelrod & Hamilton (1981), Nowak & Sigmund (1992), or Miller (1996), seem to suggest that Tit-for-tat (and other similar strategies) prevail in most scenarios. Thus, the interest of our research is putting these two approaches, mathematical and computational, side-by-side to achieve a better understanding of the evolutionary behavior of players in the repeated prisoners’ dilemma

2 The Mathematical Model

The basic stage game (Prisoners’ Dilemma) that players will play repeatedly is given by

	C	D
C	3,3	0,5
D	5,0	1,1

We now consider the repeated version of the game played a finite number of rounds R . In order to have a tractable model, we only consider three possible strategies (as in Imhof *et al.* (2005)):

- D : Always defect
- C : Always cooperate
- T : Tit-for-Tat

as they are the three strategies that have deserved a higher attention in almost all the literature dealing with the Repeated Prisoner's Dilemma from an evolutionary point of view. The fact that we only consider 3 possible strategies clearly imposes a strong restriction to the analysis, as we will discuss later. Given the above, the repeated game can be represented as follows

	C	D	T
C	$3R, 3R$	$0, 5R$	$3R, 3R$
D	$5R, 0$	R, R	$5 + (R - 1), (R - 1)$
T	$3R, 3R$	$(R - 1), 5 + (R - 1)$	$3R, 3R$

Thus, for instance, when a D -type strategy meets a T -type strategy, the former gets 5 in the first round and then 1 in each subsequent round ($5 + (R - 1)$ in total), while the later gets 0 first and then 1 in each subsequent round ($R - 1$ in total).

2.1 The Replicator Dynamics analysis

Let $p_t(C)$ be the probability that, at time t , a player in this population is an “always cooperate” type, and the same for $p_t(D)$ and $p_t(T)$. We thus have that $p_t(C) + p_t(D) + p_t(T) = 1 \forall t$.

The replicator dynamics states that the rate of change of such probabilities is a function of the relative performance of each strategy with respect to the average performance of the population. In this sense, given $p_t(C)$, $p_t(D)$, $p_t(T)$, the expected payoff at time t for each strategy ($E_t \pi(\cdot)$) is:

$$E_t \pi(C) = 3R p_t(C) + 0 p_t(D) + 3R p_t(T) = 3R(p_t(C) + p_t(T))$$

$$E_t \pi(D) = 5R p_t(C) + R p_t(D) + (5 + (R - 1)) p_t(T)$$

$$E_t \pi(T) = 3R p_t(C) + (R - 1) p_t(D) + 3R p_t(T) = 3R(p_t(C) + p_t(T)) + (R - 1) p_t(D)$$

and thus the average payoff will be:

$$E_t \bar{\pi} = E_t \pi(C) p_t(C) + E_t \pi(D) p_t(D) + E_t \pi(T) p_t(T)$$

Notice that since $p_t(C) + p_t(D) + p_t(T) = 1 \forall t$ only two dimensions matter. Hence, the replicator dynamics in this case is given by:

$$\frac{\partial p_t(C)}{\partial t} = p_t(C)(E_t \pi(C) - E_t \bar{\pi})$$

$$\frac{\partial p_t(D)}{\partial t} = p_t(D)(E_t \pi(D) - E_t \bar{\pi})$$

The corresponding vector field showing the trajectories of the system is represented in Figure 1, where the points a and b given by

$$a = \frac{2R-4}{2R-3}, \quad b = \frac{R-2}{R-1}$$

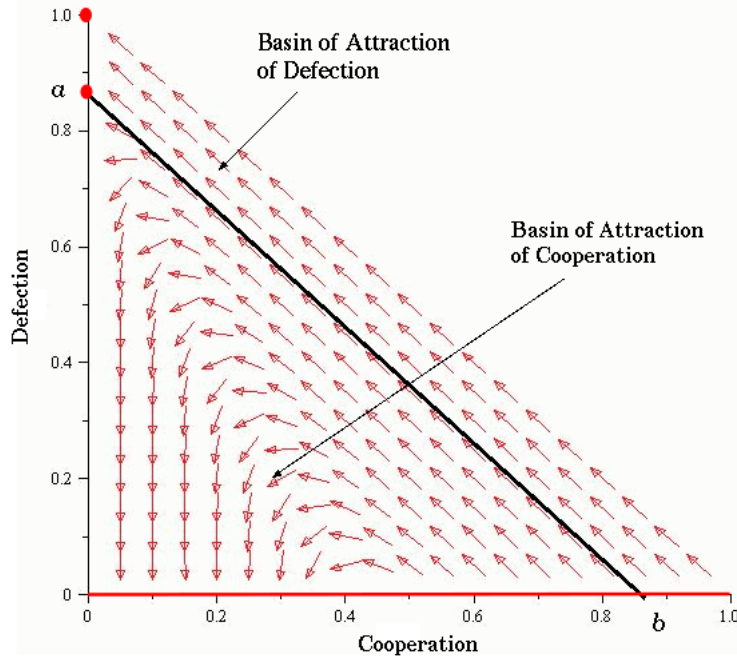


Fig. 1: Vector Field

The horizontal and vertical axis in Figure 1 correspond to $p_t(C)$ and $p_t(D)$ respectively. Thus, the three vertexes of the triangle $((1,0)$, $(0,1)$, and $(0,0)$), correspond to the states $p_t(C) = 1$, $p_t(D) = 1$, and $p_t(T) = 1$ respectively.

The trajectories that represent the evolution of the system are divided in two areas or *basins of attraction*, one for Defection and another for Cooperation. Thus, depending on the location in the simplex of the initial probabilities at $t = 0$, $(p_0(C), p_0(D))$, the system will evolve according to the corresponding path towards defection (the vertex $(0,1)$) or cooperation (somewhere along the line $(0,0) \rightarrow (1,0)$). It is clear that the *basin of attraction of cooperation* grows as R becomes large. Indeed, we have that $a \rightarrow 1$ and $b \rightarrow 1$ as $R \rightarrow \infty$

Stationary points (rest points) of the system are marked red:

- $(0, 1)$, that corresponds to everybody playing *always defect* $p_t(D) = 1$,
- $(0, a)$, which is a singular point,
- all the points in the line that goes from $(0,0)$ to $(1,0)$ that correspond to points with no *defectants*, that is, $p_t(D) = 0$ and $p_t(C) + p_t(T) = 1$.

Notice that only the point $(0, 1)$ corresponding to $p_t(D) = 1$ is asymptotically stable in the sense that if the system is slightly perturbed away from $(0, 1)$, any trajectory will bring it back to the same point. The singular point $(0, a)$, which is not asymptotically stable, can only be reached if the system starts somewhere in the line that goes from $(0, a)$ to $(b, 0)$, which occurs with zero probability. Finally, points in the line $(0, 0) \rightarrow (1, 0)$ are stationary but not stable.

An important result is that the relative size of these basins of attraction depends on the number of repetition R . That is, if the system starts at random, the probability of reaching the point $(0, 1)$ ($p(0, 1)$, everybody defecting) or the line $(0, 0) \rightarrow (1, 0)$ ($p((0, 0) \rightarrow (1, 0))$, everybody cooperating) depends on R .

$$p((0, 0) \rightarrow (1, 0)) = \left(\frac{2R-4}{2R-3}\right)\left(\frac{R-2}{R-1}\right)$$

$$p(0, 1) = 1 - \left(\frac{2R-4}{2R-3}\right)\left(\frac{R-2}{R-1}\right)$$

Thus, we can compute the *expected per-round payoff* ($E\bar{\pi}$) as a function of R .

$$E\bar{\pi} = \left(\frac{2R-4}{2R-3}\right)\left(\frac{R-2}{R-1}\right) \cdot 3 + \left(1 - \left(\frac{2R-4}{2R-3}\right)\left(\frac{R-2}{R-1}\right)\right) \cdot 1$$

Figure 2 shows the behavior of such *expected per-round payoff* as a function of R . We observe that it grows rapidly as the number of repetitions (R) increases. In fact, $E\bar{\pi} \rightarrow 3$ as $R \rightarrow \infty$

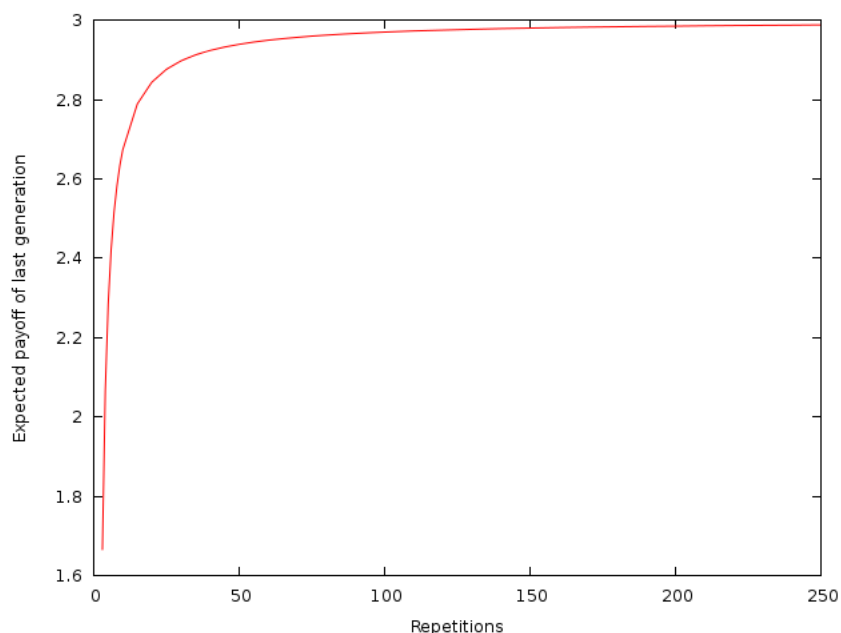


Fig. 2: Expected Payoff as a function of R

3 The Computational Model

Given the analysis above, the dynamics seem to suggest that there is room for cooperation. At least for a broad range of initial conditions, the trajectories lead to some point in the horizontal axis corresponding to a population consisting of *only* C and T strategies.

Nevertheless, such analysis is extremely partial since we are only considering 3 strategies at a time, namely C , D , and T . One can easily see that extending this approach (mathematical) to a more general case (with more strategies considered) is a difficult task as it would be extremely difficult to study the behavior of a dynamic system with more than 2 dimension

To overcome this limitation, we develop a computer simulation¹ in which the strategies are represented by finite automata of size four (encoded as binary strings of 0's and 1's) and a Genetic Algorithm routine is used to simulate the evolutionary process as in Miller (1996). The algorithm was run for 5000 generations starting from an initial random population of 100 strategies using the standard single-cut crossover operator and with a probability of per-bit mutation of 0.005. In most of the cases, the results of such simulations produce the outcome in Figure 3, in which the evolution of the (per round) average payoff is displayed.

¹ The software used for this simulation consists of a set of routines written in ANSI C. It is available from <http://ideas.repec.org/c/aub/grecss/001.05.html>

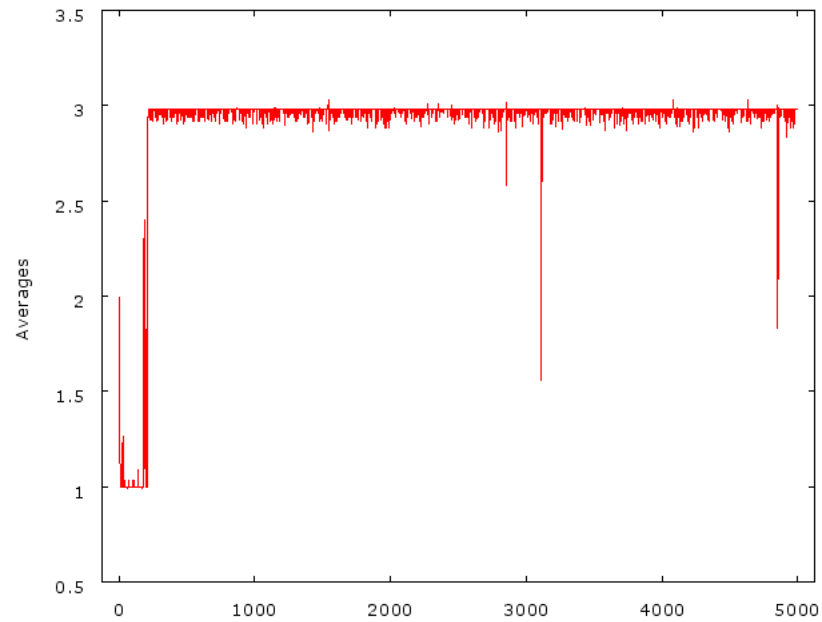


Fig. 3: Evolution of the average per-round payoff when Cooperation is the result

Because the final average payoff is 3 we can conclude that all players follow a cooperative strategy.

In other cases, though, cooperation is not the final result as the evolution of the average payoff results as in Figure 4, which corresponds to the case of all the players defecting.

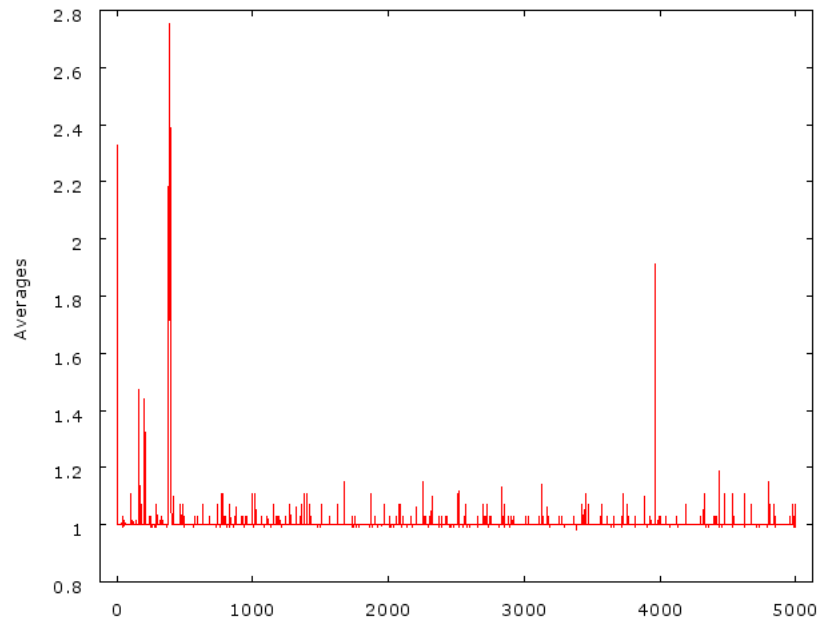


Fig. 4: Evolution of the average per-round payoff when Defection is the result

In both cases, though, the resemblance between the vector field in Figure 1 and the evolution of payoffs in Figures 3 and 4 is very appealing:

- When the final result is *cooperation* (as in Figure 3), both in the replicator dynamic analysis and in the simulations, the evolutionary process seems to favor the growth of Defectant strategies at first, and then these disappear and Cooperative strategies start to replicate to end up with the payoff corresponding to the cooperative behavior.
- On the contrary, when the final result is *defection* (as in Figure 4), the evolution goes “monotonically” towards that point.

How often each of these two results occurs in the simulations ? Given that in the mathematical model we have found that the answer to this question depends upon the number of repetitions R , we check whether R also has an effect in the computational model. In this sense, Figure 5 complements Figure 2 by showing how the final observed average payoff of the simulations² (the payoff of generation 5000) depends on R . For robustness, we perform such exercise with two different crossover rules (the “canonical” single-cut crossover and a fifty-fifty crossover³) and with no crossover

² For each value of R we run 1000 simulations and compute the average payoff of the last generation (5000)

³ Fifty-fifty crossover consists of generating new binary strings in such a way that each “locus” of the new string has 0.5 probability of being a copy of the corresponding *locus* of the “father” string and 0.5 probability of that of the “mother’s” string

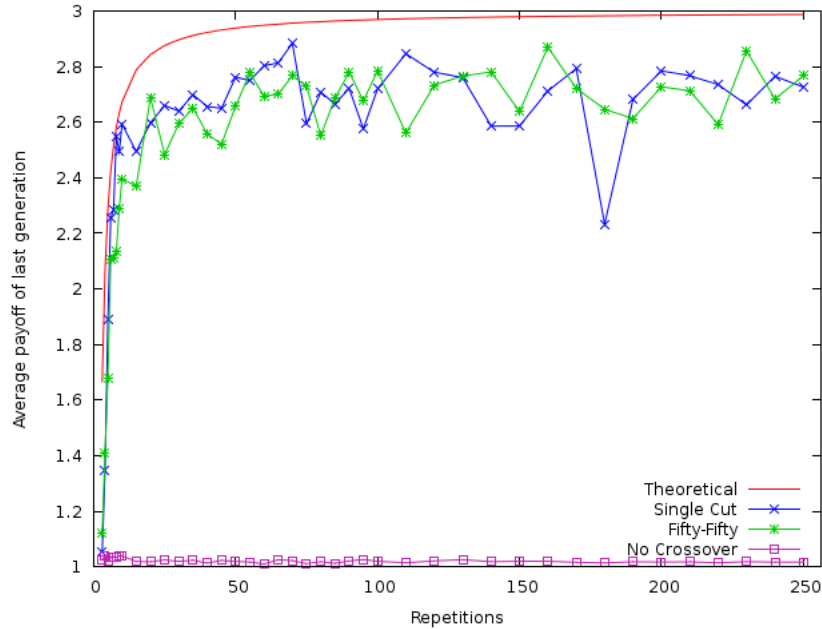


Fig. 5: Expected and observed payoffs as a function of R

Figure 5 shows to what extent the behavior of the simulations (curves marked “x” and “*”) resembles what we obtained mathematically in the previous section (“Theoretical” curve). We observe that, as the number of repetitions R grows, the higher is the probability of reaching cooperation at the end and hence, the higher is the average payoff, both theoretical and empirically.

In this sense, it seems that the use of Genetic Algorithms to simulate the evolutionary process closely matches the behavior predicted by the replicator dynamics while avoiding the strong limitation of considering only 3 possible strategies.

To test “how close” these results (mathematical and computational) are, Figures 6 and 7 show how statistically significant is the hypothesis that the mean of the average payoffs of the computational model equals the theoretical expected payoff of the mathematical model. To this purpose the dotted line corresponds to the lower end of a 95% confidence interval

We observe that, specially in the case of the “canonical” single-cut crossover and except for a few atypical observations, the results are significantly close.

4 Conclusions

We have studied the evolution of strategies in the well known Repeated Prisoner’s Dilemma using two different approaches: one *mathematical* based on the replicator dynamics and one *computational* based on genetic algorithms. We show that the results obtained from the two approaches coincide to a great extent in the sense that,

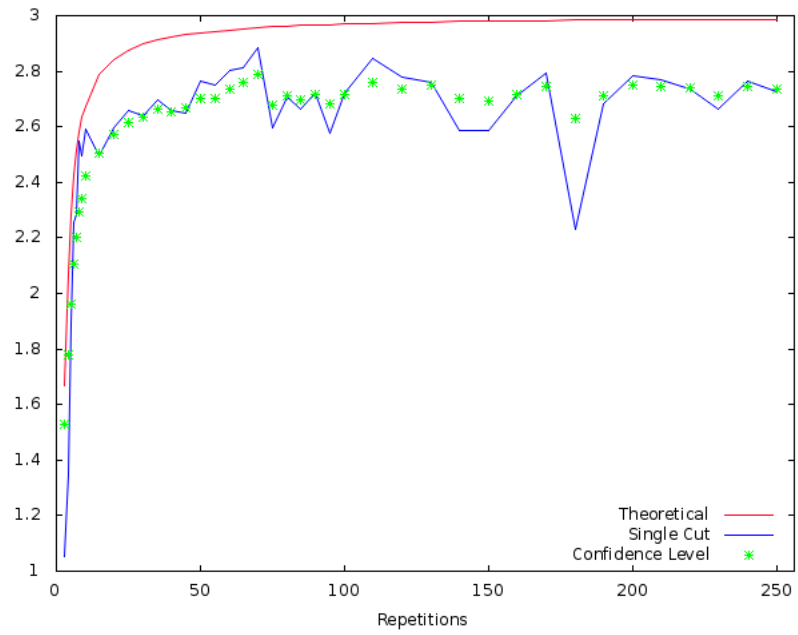


Fig. 6: Confidence analysis for the case of the *Single Cut* crossover

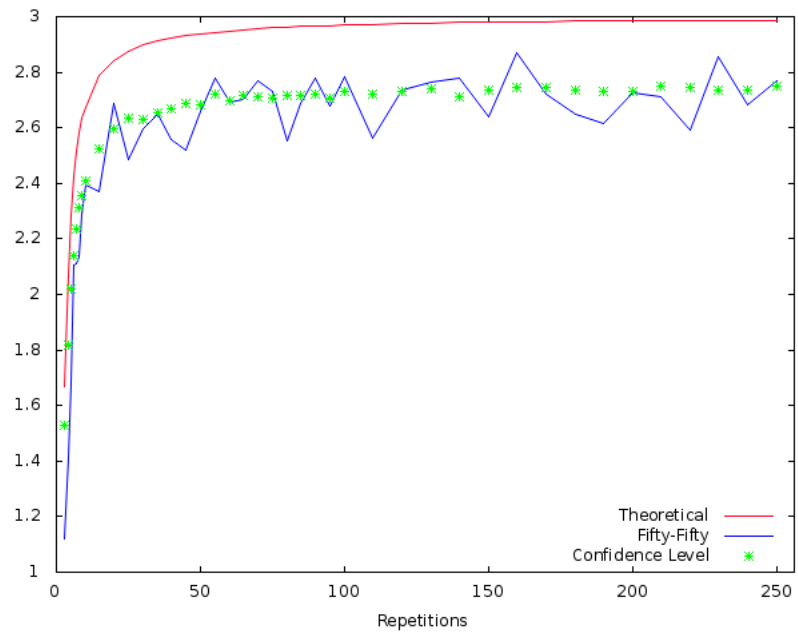


Fig. 7: Confidence analysis for the case of the *Fifty-Fifty* crossover

1. The two approaches produce the same two possible outcomes: evolution towards defection or evolution towards cooperation.
2. In the two approaches, the path towards the equilibrium are similar: monotonic when going towards defection, non-monotonic when going towards cooperation
3. In the two approaches, the percentage of times each of the two possible results occurs is also similar and depends on the number of repetitions of the game.

The scope of these conclusions, though, might be somehow limited. The reason is that, while in the mathematical model only three strategies are considered (Always cooperate, Always defect, and Tit-for-tat), the computational model deals with finite automata of size 4, which can represent a very large number of different strategies. Nevertheless, one generally observes that from a starting random population of strategies (represented by finite automata), the genetic algorithm rapidly reduces the number of “working” strategies and, at the end, only strategies similar to the three used in the mathematical model appear. Also, in Vilà (2008) we discuss other genetic algorithm operators that can deal with this issue, and the results are not different from the ones presented here. Another limitation of the present analysis is that it focuses exclusively on the Repeated Prisoner’s Dilemma. On the one hand this might be considered a good research strategy as such game has been extensively analyzed and it is very easy to put in contrast the results obtained here with other results in the literature. On the other hand the same technique should be tested with other games to verify the conclusions drawn in this paper, which is the topic of future research. Nevertheless, in Vilà (2008) a similar approach (combining mathematical and computational models) has been used to study a model of Bertrand competition with strategic sellers and buyers and the results there corroborate the main findings here: the outcomes from the replicator dynamics model and from the genetic algorithms model coincide to a high degree.

The results of this research seem to suggest that, in our opinion, mathematical and computational models are good complements for research in social sciences. Indeed, while on the one hand computational models are extremely useful to extend the scope of the analysis to complex scenarios hard to analyze mathematically, on the other hand formal models can be extremely useful to verify and to explain the outcomes of computational models without the need of resorting to verbal or *ad-hoc* explanations. For instance, in the particular example studied in this paper, a researcher looking only at the results of the computational model might conclude that the model is somehow inconclusive. Indeed, for the same set of initial conditions the simulation may end in “all cooperating” or “all defecting”, as we have seen in section 3 (Figures 3 and 4). Only the knowledge of a formal mathematical model that is aligned with the computational model can rigorously explain this apparently puzzling result.

References

AXELROD, R. (1985). *The Evolution of Cooperation*. New York: Basic Books.

- AXELROD, R. (1997). Advancing the art of simulation in the social sciences. *Complexity* **3**(2), 16–22.
- AXELROD, R. & HAMILTON, W. (1981). The evolution of cooperation. *Science* **211**(4489), 1390–1396.
- BINMORE, K. & SAMUELSON, L. (1992). Evolutionary stability in repeated games played by finite automata. *Journal of Economic Theory* **57**(2), 278–305.
- BOYD, R. & LORBERBAUM, J. P. (1987). No pure strategy is evolutionarily stable in the repeated Prisoner’s Dilemma game. *Nature* **327**(6117), 58–59.
- HEGSELMANN, R. & WILL, O. (2008). A Replication That Failed on the Computational Model in ‘Michael W. Macy and Yoshimichi Sato: Trust, Cooperation and Market Formation in the U.S. and Japan. Proceedings of the National Academy of Sciences, May 2002’. *Journal of Artificial Societies and Social Simulation* **11**(3)3. URL <http://jasss.soc.surrey.ac.uk/11/3/3.html>.
- IMHOF, L. A., FUDENBERG, D. & NOWAK, M. A. (2005). Evolutionary cycles of cooperation and defection. *Proceedings of the National Academy of Sciences* **102**(31), 10797–10800. URL <http://www.pnas.org/content/102/31/10797.short>.
- IZQUIERDO, L. R., IZQUIERDO, S. S., GALÁN, J. M. & SANTOS, J. I. (2009). Techniques to Understand Computer Simulations: Markov Chain Analysis. *Journal of Artificial Societies and Social Simulation* **12**(1)16. URL <http://jasss.soc.surrey.ac.uk/12/1/6.html>.
- MACY, M. & SATO, Y. (2010). The Surprising Success of a Replication That Failed. *Journal of Artificial Societies and Social Simulation* **13**(2)9. URL <http://jasss.soc.surrey.ac.uk/13/2/9.html>.
- MILLER, J. (1996). The coevolution of automata in the repeated Prisoner’s Dilemma. *Journal of Economic Behavior & Organization* **29**(1), 87–112.
- NOWAK, M. A. & SIGMUND, K. (1992). Tit for tat in heterogeneous populations. *Nature* **355**(6357), 250–253. URL 0.
- VILÀ, X. (2008). A Model-To-Model Analysis of Bertrand Competition. *Journal of Artificial Societies and Social Simulation* **11**(2)11. URL <http://jasss.soc.surrey.ac.uk/11/2/11.html>.
- WILL, O. (2009). Resolving a Replication That Failed: News on the Macy and Sato Model. *Journal of Artificial Societies and Social Simulation* **12**(4)11. URL <http://jasss.soc.surrey.ac.uk/12/4/11.html>.