

Stochastic approximation to understand simple simulation models

Segismundo S. Izquierdo ^{a,b} and Luis R. Izquierdo ^{b,c}

^a*Department of Industrial Organization, EII, Universidad de Valladolid, 47011 Spain.*

^b*InSiSoc, Social Systems Engineering Centre, P del Cauce 59, 47011 Valladolid, Spain.*

^c*Department of Civil Engineering, Universidad de Burgos, 09001 Spain.*

Abstract

This paper illustrates how a deterministic approximation of a stochastic process can be usefully applied to analyse the dynamics of many simple simulation models. To demonstrate the type of results that can be obtained using this approximation, we present two illustrative examples which are meant to serve as methodological references for researchers exploring this area. Finally, we prove some convergence results for simulations of a family of evolutionary games, namely, intra-population imitation models in n -player games with arbitrary payoffs.

Keywords: stochastic approximation; mean dynamic; Markov models; evolutionary games.

1 Introduction

This paper is intended to be useful for researchers that are interested in the relation between a computer simulation of a discrete-time stochastic process and the expected move of such stochastic process, also called its mean dynamic or mean-field deterministic approximation. Specifically, we show how stochastic approximation theory can help to understand the dynamics of various computer simulations. Here we focus on models that can be formulated as Markov chains, complementing other more basic Markov-based techniques for the analysis of computer simulations [1, 21, 28]. The principle underlying the use of a mean-field approximation to analyse the behaviour of a system is to develop a simplified model based on the average or expected value of some of the variables or magnitudes in the system, ignoring or partially ignoring the fluctuations of that magnitude around its average. Under the umbrella of such a broad principle, or even when considering slightly more restrictive frameworks, there are many different mean-field approximation techniques [39]. The particular one that we study in this paper derives from Stochastic Approximation Theory (SAT), and presents

some differences with the technique derived from Statistical Physics known as Mean Field Theory (MFT). For the sake of clarity, we devote here a few lines to clarify the main differences between these two approaches.

Mean-field approximations were introduced in physics in the early 1900s through the work of Weiss on magnetism [49]. This technique approximates the simultaneous effect –or field– that a number of interacting neighbouring particles of different types (or in different states) create at a given location, by considering some average value or mean field which may depend only on the total number of neighbouring particles at the site and on the global –rather than local– prevalence of each type of particle.

This kind of mean-field approximation has proved to be very useful in a large number of fields, from physics to sociology [19, 26], and has given rise to an enormous number of variants and improvements, including dynamical extensions –especially to study strongly correlated electronic systems [31]. In particular, the technique has been usefully applied to study dynamics in networks [2] and evolutionary games on graphs (see [46] and [42] for a review).

The starting point of the mean-field approach in Stochastic Approximation Theory is very different. SAT studies recursive algorithms such as $X_{n+1}^\gamma = X_n^\gamma + \gamma \cdot Y_n^\gamma$, where γ is a positive real parameter which can be –or become– small, and where Y_n^γ is a vector of random variables (i.e. a stochastic term) whose distribution may depend on the value of X_n^γ . In this framework, “mean-field approximation” or “mean dynamic”[44] refers to a model in which the stochastic term Y_n^γ is approximated by its expected value $E(Y_n^\gamma)$ at any point $X_n^\gamma = x$, and in which the equation in differences $X_{n+1}^\gamma = X_n^\gamma + \gamma \cdot Y_n^\gamma$ is approximated by a deterministic differential equation $\dot{x} = f(x)$. In the context of recursive stochastic algorithms, this method was introduced in 1977 by Ljung [36], and, for stochastic learning models, it can be traced back to the work of Norman in the 1960s [40, 41]. Besides *mean-field approximation*, *mean-field equations* or *deterministic approximation* [7, 13], the method has also received several other names, such as *mean dynamic* [44], *approximate mean value equation* [46] or *the (ordinary differential equation) ODE method* [33]. In contrast to the Statistical Mechanics approach described above, in the stochastic approximation case there is no averaging of concurrent effects. Instead, averaging (or expectation) in SAT is done across time rather than space, and a large number of small *sequential* steps (small values of γ) will be needed for the

approximation to be valid. In fact, as illustrated in some of the examples we will discuss later, there may be no interacting particles at all, and the technique can be applied to a single particle subject to a random external field.

As to the relevance of the number of components or particles in a system, it is not rare that, when modelling some real systems, the step parameter γ in the recursive equation $X_{n+1}^\gamma = X_n^\gamma + \gamma \cdot Y_n^\gamma$ is a decreasing function of the number of interacting agents in the system (often in the form of inverse proportion) [6]. In those particular cases, a large number N of particles –or agents– will also be the key condition (as is usually the case in Mean Field Theory) for the SAT mean-field approximation to be valid, but, in contrast to the usual MFT case, it may be the case that only 2 particles or agents are interacting at each step. The important point for the stochastic approximation in SAT to be valid is that the “global state” of the population varies very little between consecutive time steps (i.e. many steps should take place before the “global state” can change significantly); this condition is often satisfied when the population is very large, the global state of the population can be characterised by the number of agents of each type, and only a few agents can change their state at each time step. In contrast, in MFT, the important condition is usually that a sufficiently large (but possibly small, e.g. greater than four) number of particles interact *simultaneously* with any given one [22]. The stochastic approximation will usually get gradually better as the population size N increases, and very large values of N may be needed for the approximation to be a good description of the dynamics of the system over a limited period of time (steps).

In a real system, both the SAT and MFT approaches can lead to the same set of equations for the dynamics of some average value in the population, but then the MFT approach will usually be interpreted as the dynamics of the corresponding continuously varying variable under some homogeneity assumption in an infinite population [2], while the SAT approach will be providing an approximation to the actual stochastic trajectory of the variable in a finite but large population, when that variable changes very little between consecutive steps.

From now on, in this paper we will use the name *mean dynamic* to refer to the ODE method derived in Stochastic Approximation Theory. Our approach is threefold: First, we intuitively discuss the type of models for which the mean dynamic can be a useful tool. Then we present some technical conditions that guarantee the validity of this approximation. These are selected results from the theory of stochastic

approximation, which we have appropriately simplified for the objectives of this paper. And finally, as an illustration of the possible applications of the theory, we provide examples and proofs of convergence for simulations of various 2×2 and n -player evolutionary games.

We do not provide new theorems in stochastic approximation: most of the theoretical results in this paper can be derived from a general theorem in Stochastic Approximation Theory proved by Norman [40]. In this sense, the first part of the paper is mainly an introduction to the mean dynamic as a tool to understand computer simulations, and the added value of this paper is that it selects and adapts those results that we consider most useful for understanding the dynamics of simple simulation models, and presents the corresponding technical requirements in a simpler way than the original sources.

A second goal of this paper is to provide the reader with the intuition that underlies the mean dynamic. This is valuable because –as illustrated below– the mere intuition is often sufficient to significantly enhance our understanding of the dynamics of various computer simulations. Furthermore, it turns out that the intuition is often useful even in cases where the precise mathematical theorems that back such intuition formally have not yet been developed (and the theory of stochastic approximation is a very dynamic field, so today’s intuitions may become tomorrow’s theorems). We also discuss precise technical conditions that guarantee the formal validity of the mean dynamic, but the value of comprehending the intuition should not be underestimated.

Mean dynamic approximations of stochastic models have been widely and successfully applied to analyse a diverse range of models in e.g. learning game theory [10, 18, 24, 27, 29, 48], evolutionary game theory [4, 5, 7, 9, 14, 30] and diffusion processes in social networks [19, 37, 47]. As a particular application of the theory, we show in this paper some agent-based models of social interactions which present completely different dynamics depending on the number of agents included in the model (this is already a well-known fact), and we explain how these qualitatively different behaviours can be characterised.

The rest of the paper is structured as follows: section 2 explains the difference between transient and asymptotic behaviour in the context of Markov models. In section 3 we informally discuss the conditions under which the mean dynamic can be useful to understand the transient dynamics of a computer simulation. As mentioned above, it is important to note that mean dynamic methods can shed light on the dynamics of stochastic processes even in cases where there are still no readily available convergence

theorems that can prove the formal validity of the approximation. Section 4 provides a set of formal results and sufficient conditions that guarantee the validity of the mean dynamic. These results are adapted mainly from Norman [40]. Section 5 discusses in practical terms the type of analysis that can be conducted on computer simulations using stochastic approximation theory. In particular, we provide two illustrative examples of how to analyse computer simulations of simple models combining the theory of Markov chains with the theory of stochastic approximation. At the end of this section we provide generalised results for simulations of some 2x2 and n -player evolutionary games. Finally, section 6 summarises our conclusions.

The paper also includes an Appendix which provides some notes on the relation between a *deterministic* difference equation of the form $\Delta x_n = \gamma \cdot f(x_n)$ and its corresponding differential equation $\dot{x} = f(x)$. Some readers may find these notes useful before embarking on the analysis of the *stochastic* case.

2 Markov Chains. Transient and asymptotic behaviour

This section briefly revisits some of the techniques that can be used to study Markov chains. The main objective of this section is to set the scope of applicability of the mean dynamic approximation, namely the study of the transient behaviour of systems with large transition matrices (or with an infinite state space). The reader familiar with Markov chains may wish to skip this section.

We centre our discussion on discrete-time stochastic processes that can be formalised as Markov chains, i.e. processes in which, in every time-step n , a vector of state variables X_n can condense all the past information of the process that is relevant for its future. In a way, the values of such variables depict a “complete and accurate snapshot” of the current situation of the system. Thus, the value of the vector of state variables X_n is often called “the state of the system in time-step n ”. The current state of the system, together with the state transition probabilities, determine the probability distribution over the possible states of the system for every future time-step. Slightly more formally, given some initial state $X_0 = i$, we can calculate the probability distribution of the state of the system after n time steps: $P(X_n = \cdot \mid X_0 = i)$.

The study of Markov processes is usually divided into transient and asymptotic behaviour. The transient behaviour is characterised by the probability distribution of the

state vector X_n in a given time-step n . The asymptotic behaviour is characterised by the limit of the distribution of X_n as n goes to infinity, when this limit exists.

After a sufficiently long period of time, a Markov process –and, consequently, a computer simulation of the process– is likely to be evolving very closely to what its asymptotic behaviour (stochastically) predicts; as a result, many studies of Markov processes focus on their asymptotic dynamics. Techniques to characterise the asymptotic behaviour of Markov chains can be found in introductory textbooks [32]. However, “a sufficiently long period of time” may be too long, i.e. it may be unattainable in practical terms (e.g. it may require years of computation) or, simply, it may be significantly longer than the time scope we are actually interested in. In many cases, as it happens so often in life, the journey may be more interesting than the final state.

Thus, let us now turn to the transient behaviour. Consider a Markov chain with s possible states. In simple words, we are after an s -dimensional vector $a^{(n)} = [a_1^{(n)}, \dots, a_s^{(n)}]$ containing the probability of finding the process in each of the s possible states in time-step n . The i^{th} element of $a^{(n)}$ is $a_i^{(n)} = P(X_n = i)$, and it denotes the probability that the system is in state i at time-step n . To calculate $a^{(n)}$ we must first introduce the transition matrix \mathbf{P} . Let $p_{i,j}$ denote the probability that the system moves from state i to state j in one time-step: $p_{i,j} = P(X_{n+1} = j | X_n = i)$. The probabilities $p_{i,j}$ are called transition probabilities and they are often arranged in a matrix, namely the transition matrix \mathbf{P} . This matrix \mathbf{P} characterises the stochastic process.

If the transition matrix \mathbf{P} of a Markov chain is known and tractable, computing the transient behaviour of the system is straightforward: the probability distribution of the state of the system in time-step n is characterised by $a^{(n)} = a^{(0)} \cdot \mathbf{P}^n$. (Naturally, in general the distribution of the state of the system in time-step n depends on the initial conditions $a^{(0)}$.)

If the state transition matrix is either unknown or intractable, the situation is not so simple. In some cases, given some initial conditions $a^{(0)}$, computer simulations can approximate the distribution of the state of the system in time-step n , and the statistical accuracy of these approximations can be assessed [28]. Unfortunately, the nature of the parameter space of many models means that obtaining good approximations for their transient behaviour in the general case (i.e. for any parameterisation and all initial conditions) is often too demanding in computational terms. When this is the case, can we still say something about the expected behaviour of the Markov process before it

gets close to its asymptotic behaviour? This paper shows that sometimes we can, by using stochastic approximation theory. The following sections explain how the mean dynamic can be useful to understand the transient dynamics of some simple simulation models formulated as Markov chains.

3 The mean dynamic: Intuition

A simulated example can illustrate the conditions under which the mean dynamic can be useful to understand a computer simulation. Consider one agent situated at a point in a bounded 2-dimensional space. In each time-step this agent takes one step in one of the four cardinal directions. The agent's movement is stochastic and slightly biased towards the North: a northward step is twice more likely than a step in any of the other three cardinal directions. Suppose that we place several agents like this one in the middle of a small area, such as a garden, and several others in the middle of a very large area, such as a desert (see applet for Fig. 1 at <http://luis.izqui.org/models/gardendesert/>).

If we observe the behaviour of the agents in the garden before they leave it, it will probably look quite random, with the agents leaving the garden through all its different sides. In contrast, if we observe the behaviour of the agents in the desert, they will seem quite determined to advance northwards, and with high probability they will all be leaving the desert through its northern border. When the number of steps required to cover some distance is very large, the randomness of the agents' movement averages out, and the agents will seem to be following their expected move northwards. This is the basic idea of the mean dynamic.

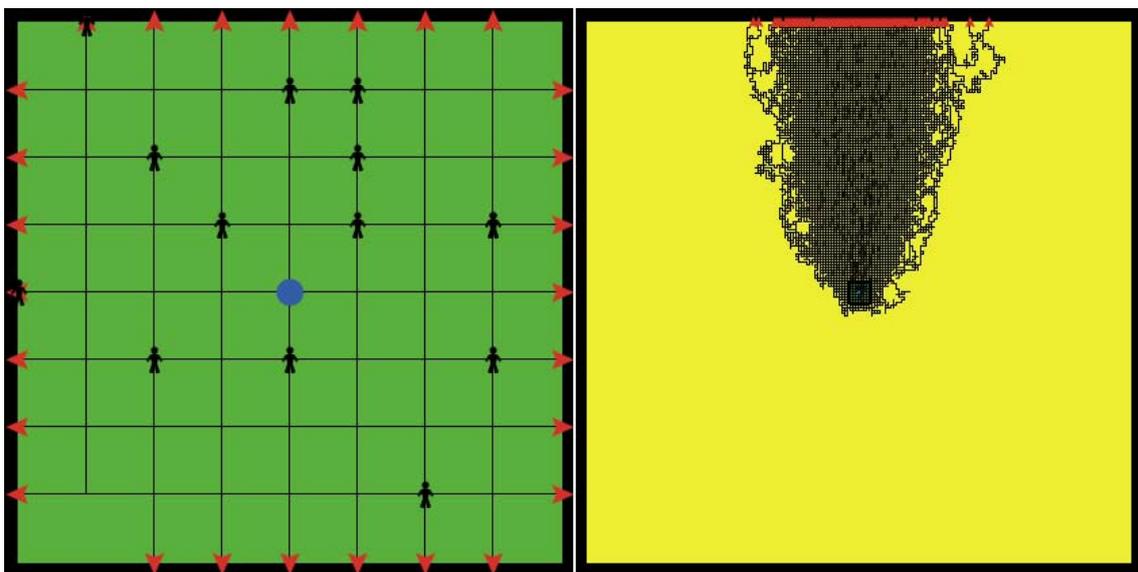


Figure 1. The applet shows two independent areas: a small (8x8) green garden on the left, and a large

(200x200) yellow desert on the right. For each area, an iteration of the model implies (a) the creation of one agent in the middle of the area, and (b) that each agent within the area advances one step: upwards (with probability 0.4), leftwards (with probability 0.2), rightwards (with probability 0.2) or downwards (with probability 0.2). Agents that reach the border of their corresponding area cannot return. The button “Clear” initialises the model. Applet at <http://luis.izqui.org/models/gardendesert/>.

Loosely speaking, a requisite for the mean dynamic to be useful is that any two consecutive states of the process are “very close” to each other –according to some measure of distance in the space of states– so the process needs to take a large number of time-steps in order to advance a significant distance in that space. Simply put, time proximity must imply space proximity. Consequently, we need to work with some measure of distance between different states. As in the previous example, we will assume that the state variables of the stochastic process –the agents’ position in our example– take real values within a closed and bounded subset of d -dimensional Euclidean space \mathbb{R}^d .

Going back to our northward-biased random walker, suppose that the length of his step is γ , and let the 2-dimensional vector X_n^γ be his position vector in the grid at time n .

Note that the law of motion of an agent with step size γ is $\Delta X_n^\gamma = \gamma \cdot Y_n$, where Y_n is a random vector which may take the value $[0, 1]$ (northwards) with probability 0.4 or the values $[1, 0]$, $[-1, 0]$ or $[0, -1]$ with probability 0.2 each. Note also that the expected move of X_n^γ is $E(\Delta X_n^\gamma) = \gamma \cdot E(Y_n) = \gamma [0, 0.2]$ (i.e. northwards). The mean dynamic relates the actual (stochastic) move of X_n^γ with its (deterministic) expected move, for low values of the step size γ .

In the previous example the probability distribution of the random “jumps” Y_n is the same at every position X_n^γ and for every step size γ . More generally, consider a family of time-homogeneous *Markov* processes X_n^γ which take values within d -dimensional Euclidean space \mathbb{R}^d and such that $\Delta X_n^\gamma = \gamma \cdot Y_n^\gamma$, where Y_n^γ is a vector of random variables whose distribution may depend on the particular value of $X_n^\gamma = x$ and on γ , i.e. $Y_n^\gamma = Y_n^\gamma(\gamma, x)$. Let the mean dynamic for X_n^γ be the (*deterministic*) differential equation $\dot{x} = g(x)$, where $g(x) = \lim_{\gamma \rightarrow 0} E(Y_n^\gamma | X_n^\gamma = x)$. Informally, the vector field $g(x)$ provides the expected direction of the “jump” of X_n^γ at the particular but generic

point $X_n^\gamma = x$ when the step size γ goes to zero: the local direction of the expected move of X_n^γ when γ is small. The question then is: under what conditions can one legitimately expect that the original *stochastic* process will tend to move roughly in the local direction of its *expected* (deterministic) move?

Intuitively, if the vector field $g(x)$ associated to the local expected move of the stochastic process varies smoothly with x and the process X_n^γ is such that it takes a large number of steps for X_n^γ to leave a neighbourhood of any given point x , then we could expect that stochasticity somewhat averages out, and that X_n^γ will tend to advance in the local direction of its expected move, closely following a trajectory of the differential equation $\dot{x} = g(x)$.

Thus, the mean dynamic can be helpful to understand the dynamics of a simulation of a Markov process when it is not possible or practical to work with the state transition matrix directly and:

- i) The considered process is such that the distance between any two consecutive states is small, or decreases as the process advances. More generally, the probability that the state of the process makes a large jump to a distant area of the state space in a small number of steps must be (or become) negligible.
- ii) The process is such that, if state x' can be reached from state x in a small number of steps, then the vectors $g(x)$ and $g(x')$ that characterise the expected move of the system at those states are also very close.

These two conditions are basically “hints” or “signals” that suggest that the mean dynamic will probably be useful to understand the dynamics of the process under investigation. Some examples of Markov processes where the conditions required for the mean dynamic to be useful are usually met are:

- Dynamic models of large populations where in every time-step at most k individuals may change their strategy (i.e. their behaviour) within a finite set of strategies. This change of strategy could happen by death and birth, imitation, or other adaptation processes. If the evolution of the process depends only on the proportion of agents in the population following each possible strategy i , i.e. if those proportions fully characterise the state of the system, then the distance between two consecutive states gets small for large population sizes N , and therefore the conditions for the usefulness of the mean dynamic are satisfied. (Formally, let vector X_n , with components $X_{i,n}$, denote

the proportion of individuals following strategy i in time-step n . Then: $\Delta X_{i,n} \leq k/N$.)

Furthermore, the conditions may also be met if the number of individuals that may change their behaviour in time-step n follows a distribution with finite mean and sufficiently light tails [34].

- Learning processes in social interactions where the parameters that regulate an agent's behaviour vary slowly –the so-called “slow learning” [5, 40]. This may happen because the learning model assumes that experienced agents react less and less to new observations as time goes by –as it happens in *stochastic fictitious play* [17, 18, 24] and in Erev and Roth's reinforcement learning model [15, 25]–, or simply because the agents are assumed to be changing their parameters by small steps, as in the Bush-Mosteller reinforcement learning model [10, 27, 29]. In this latter case the dynamics of the state of the process can be expressed as $\Delta X_n^\gamma = \gamma \cdot Y_n$, where the distribution of Y_n depends only on the value of X_n^γ , and where γ is a constant and small step size.

4 Formal results

The propositions in this section are adapted from Norman [40, 41], who analyses Markov processes that advance by small steps. This is actually a particular case within the more general theory of stochastic approximation [34]. The exposition we present here is intended to be easier to understand than the original sources: it can be considered a corollary of Norman [40] that focuses only on selected results and tries to present them together with the underlying intuition, and with an emphasis on those aspects that can be useful to interpret the dynamics of computer simulations.

We are considering families of Markov processes whose state vector X_n^γ takes values within d -dimensional Euclidean space \mathbb{R}^d and evolve according to the relation

$\Delta X_n^\gamma = \gamma \cdot Y_n^\gamma$, where:

- $\gamma \in (0, \gamma_{Max}]$ is a positive real number that regulates the distance between consecutive states. We focus on the dynamics of these processes for small values of γ (i.e. small steps).
- Y_n^γ is a vector of random variables that usually corresponds to some observed variables of the system and whose distribution may depend on the particular values of X_n^γ and γ .

For any (row) vector y let $|y| = \sqrt{y \cdot y^{Tr}}$, where Tr indicates transposition, and for any $N \times N$ matrix $A = [a_{i,j}]_{N \times N}$ let $|A| = \sqrt{\sum_{i,j=1}^N a_{i,j}^2}$.

The required conditions are:

Assumption 1

Intuitively, this condition states that the transition probabilities of the Markov process do not depend on the time-step index n (stationarity), that the state of the process takes values within a convex set I of \mathbb{R}^d and that, as the step size γ goes to zero, the state space “fills” the set I , in the sense that the state of the process may be arbitrarily close to any point x in the set I .

Formally: For every γ , $\{X_n^\gamma\}_{n \geq 0}$ is a Markov process with stationary transition probabilities and state space a subset I_γ of \mathbb{R}^d . Let I be the smallest closed convex set of \mathbb{R}^d including all I_γ . We assume that I_γ approximates I as $\gamma \rightarrow 0$ in the sense that, for any $x \in I$, $\lim_{\gamma \rightarrow 0} \inf_{x_\gamma \in I_\gamma} |x_\gamma - x| = 0$.

Assumption 2

Intuitively, this condition requires the existence of the mean-field function $g(x)$, which provides the direction of the mean field at any point $x \in I$ when the step size γ goes to zero. It also requires some regularity conditions on $g(x)$.

Formally: Let $g(x, \gamma)$ be the expected value of Y_n^γ when $X_n^\gamma = x$, i.e., $g(x, \gamma) = E(Y_n^\gamma | X_n^\gamma = x)$. Because of stationarity, this function does not depend on n . We assume that:

- i) $g(x, \gamma)$ converges to a function $g(x)$ as $\gamma \rightarrow 0$ and this convergence is such that $g(x, \gamma) = g(x) + O(\gamma)$ uniformly in $x \in I_\gamma$.¹
- ii) $g(x)$ is differentiable on I and its (matrix) derivative $g'(x)$ is bounded and Lipschitz continuous on I : $\sup_{\substack{x, y \in I \\ x \neq y}} \frac{|g'(x) - g'(y)|}{|x - y|} < \infty$.

¹ A function $f(x, \gamma)$ is $O(\gamma)$ as $\gamma \rightarrow 0$ uniformly in $x \in I$ iff $\exists \delta > 0, \exists M > 0$ such that $|f(x, \gamma)| \leq M \cdot |\gamma|$ for $|\gamma| < \delta$ and for every $x \in I$, where δ, M are constants (independent of x).

Assumption 3

Let $s(x, \gamma)$ be the conditional covariance matrix of Y_n^γ when $X_n^\gamma = x$, i.e., $s(x, \gamma) = E[(Y_n^\gamma - g(x, \gamma))^{Tr} \cdot (Y_n^\gamma - g(x, \gamma)) | X_n^\gamma = x]$. We assume that, as $\gamma \rightarrow 0$, $s(x, \gamma)$ converges uniformly in $x \in I_\gamma$ to a Lipschitz continuous function $s(x)$ on I .

Assumption 4

$E(|Y_n^\gamma|^3 | X_n^\gamma = x)$ is uniformly bounded in $x \in I_\gamma$, $\gamma \in (0, \gamma_{Max}]$.

The mean limit ordinary differential equation for this kind of processes is the differential equation $\dot{x} = g(x)$. This equation plays an important role in understanding the dynamics of the process for small values of γ . Basically, with the previous conditions:

i) For any $x_0 \in I$ the ordinary differential equation (ODE) $\dot{x} = g(x)$ with initial condition $x(t=0) = x_0$ has a unique solution with trajectory $x(t, x_0) \in I$ for all $t \geq 0$.

ii) For low values of γ , the stochastic process X_n^γ with initial state $X_0^\gamma = x_0$ tends to follow the trajectory $x(t, x_0)$ of the mean limit ODE $\dot{x} = g(x)$.

More formally, for any finite t ($0 \leq t \leq T$), as $\gamma \rightarrow 0$, the stochastic process X_n^γ in time-step $n = \text{int}(t/\gamma)$ converges in probability to the point $x(t, x_0)$ of the trajectory.

iii) If, as t increases, the trajectory $x(t, x_0)$ approaches an asymptotically stable point (ASP), then, for low enough values of γ and in a sufficiently long simulation, the stochastic process X_n^γ with initial state $X_0^\gamma = x_0$ will tend to approach and linger around the ASP.

More formally, for any large enough finite $T > T_\delta$ the probability of finding the process X_n^γ in time-step $n = \text{int}(T/\gamma)$ within a small neighbourhood B_δ of the ASP goes to one as $\gamma \rightarrow 0$.

iv) For any fixed value of γ , the process X_n^γ will eventually approach its asymptotic behaviour, which –as will be shown later– cannot be characterised by the mean limit ODE in the general case [3, 14].

The results of Norman [40, 41] on learning by small steps can be considered a particular case of stochastic approximation theory. In a more restrictive framework (i.e. games with a finite number of strategies where at most one player may change her strategy

simultaneously), Benaim and Weibull [7] provide sharp results for the relationship between some evolutionary game theory Markov models and their mean dynamics. Convergence theorems for Markov processes that take values on the integer multidimensional lattice can also be found in Rozonoer [43]. General results, still within Markov processes, are provided by Ethier and Kurtz [16]. Kushner and Yin [34] present an even more general framework for the topic of stochastic approximation, but when the assumptions are relaxed the results may not be as sharp. Two other very useful general references for stochastic approximation are Benveniste, Priouret and Métivier [8], and Borkar [11]. Finally, it is worth pointing out that the theory of stochastic approximation also encompasses the analysis of processes with non-constant (decreasing) values of γ , while here we have focused on processes with constant and small γ .

5 Case studies

This section presents an overview of how to usefully apply the deterministic mean dynamic approximation of a stochastic process for the analysis of a system. This is stated in abstract terms first, and then illustrated with two specific examples. Finally, the selected theoretical results of section 4 are used to prove convergence results on simulations of a broad family of evolutionary games.

When a process meets the conditions for the mean dynamic to be useful, the dynamics of the state of the process tend to follow a trajectory of the mean limit ODE. As indicated before, in these models there is usually some distance-regulating parameter (e.g., the number of agents or some learning rate) that controls the maximum distance between any two consecutive states and, in general, the shorter this distance, the closer to a trajectory the dynamics of the system will tend to be.

Given that the transient dynamics of the model will tend to concentrate around the trajectories of the mean limit ODE, it is useful to represent the map of trajectories of the corresponding ODE, if at all possible. It is also useful to look for cycles and critical points, particularly asymptotically stable points with large domains of attraction.

Asymptotically stable points can be conceived as attractors of trajectories of the mean field. For suitable values of the distance-regulating parameter (which is usually the number of interacting elements in the simulation, or a variable denoting their sensitivity to new stimuli, or a parameter that modulates their speed of learning), asymptotically stable points act as attractors of the state of the system, and computer simulations are

likely to show dynamics that approach and spend long periods of time around these points².

Thus, the study of differential equation models of this type usually follows this approach [20, 45]:

- i) Identification of the critical points (i.e., the solutions of $\dot{x} = 0$, also called equilibria or stationary states) and limit sets;
- ii) Linear stability analysis, to study the local stability of the critical points;
- iii) Global stability analysis, to find global attractors or basins of attraction of the critical points.

Numerical simulations of the ODE can also help to analyse its behaviour. The following two cases illustrate the usefulness of the mean dynamic to understand the behaviour of some computer simulations.

5.1. An imitation model in a Hawk-Dove game

Consider a population of N individuals who, in every time-step, are randomly matched in couples for an interaction that can be modelled as a symmetric 2×2 game (N is an even integer). The possible strategies in the interaction are H (for Hawk) or D (for Dove). The preferences over outcomes for the row player are summarised in the following payoff matrix:

	H	D
H	0	3
D	1	2

where higher numbers denote higher preference. At the end of every time-step one randomly selected player revises her strategy, and her revision rule is as follows: “I look at another (randomly selected) individual in my population; if and only if I prefer her last outcome to my last outcome, I adopt her strategy”.

² See for instance the Self Correcting Equilibria (SCEs) described by Macy and Flache [38] and characterised as asymptotically stable points of the mean field by Izquierdo et al. [27].

STEP 1: Preliminary simulations

Let X_n be the fraction of the population using strategy H at the beginning of the corresponding interaction. Figure 2 shows the evolution of X_n in two simulations with 10 players, for different initial number of H-strategists in the population.

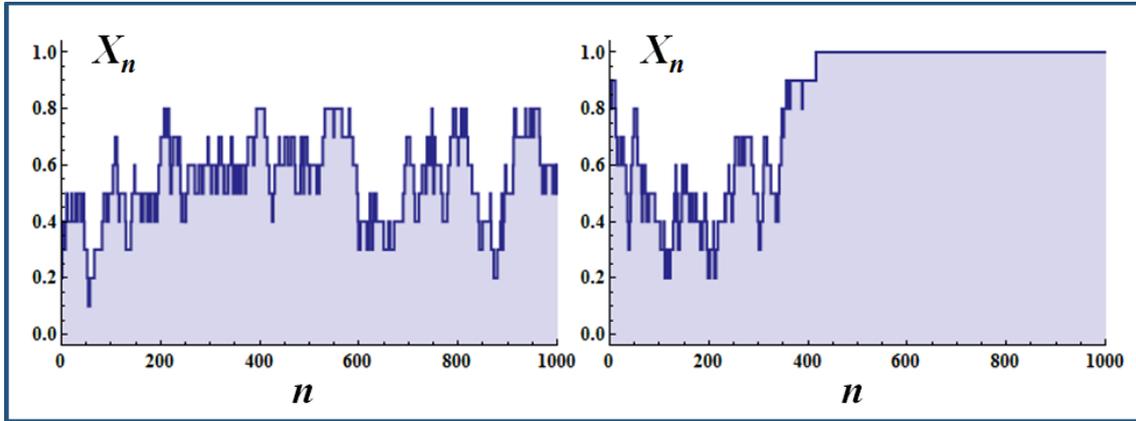


Figure 2. Evolution of X_n in two simulations with 10 players and different initial number of H-strategists in the population. Interactive figure at <http://demonstrations.wolfram.com/AnImitationModelInTheHawkDoveGame/>. Generalised version at <http://demonstrations.wolfram.com/AnImitationModelFor2x2SymmetricGames/>.

If simulations are run for long enough, we can observe that all of them end up in the (absorbing) situation $X_n = 1$. The reader can confirm this fact by running simulations with the [interactive version of figure 2](#). It is also clear in the simulations that, when the initial number of H-strategists in the population is fairly low, the system tends to spend a significant amount of time in the neighbourhood of $X_n = 0.5$.

STEP 2 Markov formulation

This model can be formalised as a Markov process whose state in time-step n is characterised by the value of X_n , i.e. the fraction of the population using strategy H at the beginning of the corresponding interaction. There are $N + 1$ possible values for the state, depending on the number of individuals using strategy H (this number ranges from 0 to N).

STEP 3 Absorbing states and limiting distribution

This imitation model has 2 absorbing states only: $X_n = 0$ and $X_n = 1$, corresponding to the situations in which everyone plays H or everyone plays D.

The preliminary simulations suggested an asymmetry between the 2 absorbing states $X_n = 1$ and $X_n = 0$. In fact, the absorbing state $X_n = 0$ is not reachable from any other state, because one single H-strategist in a population of D-strategist will obtain the maximum payoff and consequently will not change her strategy to D. Regardless of the number of individuals, the absorbing state $X_n = 1$ is reachable from every state other than $X_n = 0$, while the absorbing state $X_n = 0$ cannot be reached from any other state.

Using the standard theory of Markov chains it is straightforward to check that, unless the process starts without any H-strategist, it will necessarily end up absorbed in $X_n = 1$, i.e. sooner or later the population will be composed of H-strategists only.

STEP 4 Transient behaviour

Let us now calculate the state transition matrix for this model. Consider a population of size N and let H_n^N be the number of H-strategists in the population in time-step n .

Then, for k natural, $0 \leq k \leq N$, and given the probabilities of events defined by the revision rule of the model, we can derive the following transition probabilities:

$$P(H_{n+1}^N = k + 1 | H_n^N = k) = \frac{N-k}{N} \frac{k}{N-1} \frac{N-k}{N-1} = \frac{k(N-k)^2}{N(N-1)^2}$$

Note that the probability above, i.e. the probability that the number of H-players increases by 1, corresponds to the probability that the revising individual is a D-player who compares her result with that of an H-player matched with a D-player.

Analogously,

$$P(H_{n+1}^N = k - 1 | H_n^N = k) = \frac{k(k-1)(N-k)}{N(N-1)^2}$$

$$P(H_{n+1}^N = k | H_n^N = k) = \frac{k^2 + N(N-k-1)}{N(N-1)}$$

These transition probabilities completely determine the state transition matrix. For instance, for a population of six individuals, and numbering the states from 0 to 6 according to the number of H-strategists, the state transition matrix \mathbf{P} is

$$P = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 125/150 & 25/150 & 0 & 0 & 0 & 0 \\ 0 & 8/150 & 110/150 & 32/150 & 0 & 0 & 0 \\ 0 & 0 & 18/150 & 105/150 & 27/150 & 0 & 0 \\ 0 & 0 & 0 & 24/150 & 110/150 & 16/150 & 0 \\ 0 & 0 & 0 & 0 & 20/150 & 125/150 & 5/150 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

It is then fairly easy to calculate, given the initial conditions, the probability of finding the process in any of its possible states in any time-step. For instance, starting with half the population of H-strategists, the probability that in time-step 100 the system has reached the absorbing state made up by H-strategists only is 0.4568 for a population of 6 individuals, while this probability is 0.0062 for a population of 12 individuals. When the number of states is large, it may not be feasible to work with the state transition matrix, and it is in this case that the mean dynamic can be very useful to understand the dynamics of the system. Note that in the example we are analysing, the following relation holds:

$$\Delta X_n^\gamma = \frac{1}{N} \Delta H_n^N = \gamma \cdot Y_n^\gamma,$$

where $\gamma = N^{-1}$ is a scalar and $Y_n^\gamma = \Delta H_n^N$ is a random variable that may take the values -1 , 1 or 0 . Let us now check that this model fulfils the conditions outlined in section 4:

1.- For every N , X_n^γ takes values in $I_\gamma = \{0, \frac{1}{N}, \frac{2}{N}, \dots, \frac{N-1}{N}, 1\}$, i.e. within the closed convex set $I = [0, 1]$. The transition probabilities for X_n^γ do not depend on n . The minimum distance from any given point in $[0, 1]$ to some point in I_γ is strictly less than N^{-1} , and consequently it goes to zero as N grows.

2.- Let $x = \frac{k}{N}$ be the proportion of H-strategists in the population. From the previously calculated probabilities for ΔH_n^N , the expected value of $Y_n^\gamma = \Delta H_n^N$ when $X_n^\gamma = x$ is

$$\begin{aligned} g(x, \gamma) &= E(Y_n^\gamma | X_n^\gamma = x) = E(\Delta H_n^N | H_n^N = k) = \frac{k(N-k)(N-2k+1)}{N(N+1)^2} = \\ &= \frac{N^2}{(N-1)^2} \left[x(1-x) \left(1 - 2x + \frac{1}{N} \right) \right] = \frac{1}{1-\gamma^2} [x(1-x)(1-2x+\gamma)] \end{aligned}$$

Thus, considering that the term in square brackets is bounded on I , we have that, for $\gamma \rightarrow 0$, $g(x, \gamma) = x(1-x)(1-2x) + O(\gamma) = g(x) + O(\gamma)$ uniformly in $x \in I_\gamma$.

The function $g(x) = x(1-x)(1-2x)$ is a polynomial, and therefore it is differentiable on the compact set I and its derivative $g'(x)$ is bounded and Lipschitz on I .

3.- The function $s(x, \gamma) = E[(Y_n^\gamma - g(x, \gamma))^2 | X_n^\gamma = x]$ can be calculated from the distribution of Y_n^γ as a sum of products of polynomials, each of which uniformly converges as $\gamma \rightarrow 0$ to a polynomial on x . Consequently, $s(x)$ is a polynomial, so it is Lipschitz on I .

4.- The function $E(|Y_n^\gamma|^3 | X_n^\gamma = x)$ is uniformly bounded, since 1 is an upper bound for $|Y_n^\gamma| = |\Delta H_n^N|$ for every $x \in I_\gamma$ and for every γ .

Therefore the results described in section 4 apply.

Consider the associated limit mean ODE for X_n : the differential equation

$$\dot{x} = x(1-x)(1-2x)$$

with initial condition $x(t=0) = x_0 \in [0, 1]$.

Figure 3 shows the solution of this differential equation from time $t=0$ to time $t=10$ corresponding to an initial condition $x(t=0) = 0.2$. It also shows the results of some simulations of the discrete time process X_n^γ with $X_0^\gamma = 0.2$ for different population sizes (N) and from time-step 0 to time-step $n = 10 \cdot N$. For each of these simulations, the horizontal separation (in the t -axis) between consecutive states is N^{-1} . It can be seen that, as the population size N gets larger, i.e. as the step size $\gamma = N^{-1}$ gets smaller:

- i) The representation of the state X_n^γ of the simulated systems with initial state $X_0^\gamma = x_0$ tends to follow the solution of the ODE with initial condition $x(t=0) = x_0$.
- ii) For any real $t > 0$, the state X_n^γ of the simulated systems in time-step $n = \text{int}(t \cdot N)$ with initial state $X_0^\gamma = x_0$ tends to approach the value at time t of the solution of the ODE with initial condition $x(t=0) = x_0$.

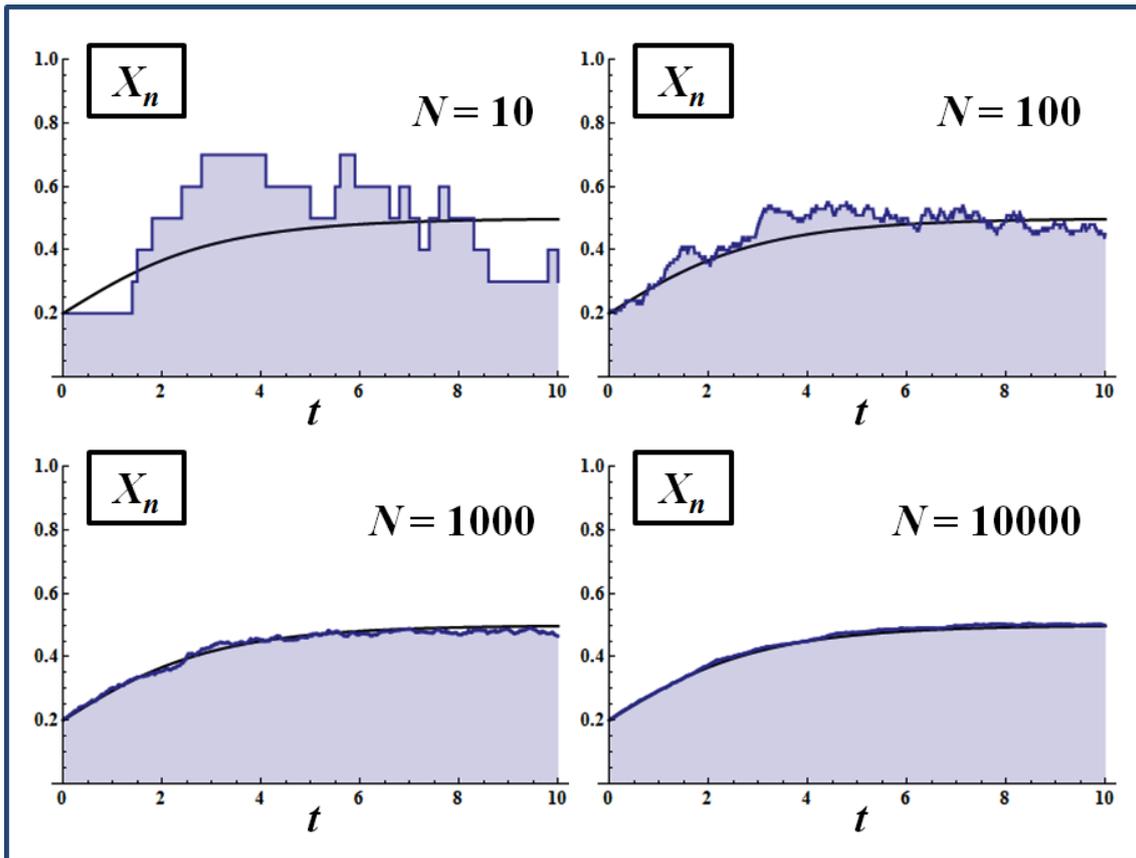


Figure 3. Evolution of X_n in four simulations with different number of players N . The initial proportion of H-strategists in the population is 0.2. The figures show the trajectory of the differential equation that describes the expected move of X_n for large values of N , which can be seen to approach $x = 0.5$. For large values of N the process tends to follow the associated trajectory of the mean field, while for small values of N , it tends to be quickly absorbed in the unique absorbing state $X_n = 1$. Interactive figure at <http://demonstrations.wolfram.com/ExpectedDynamicsOfAnImitationModelInTheHawkDoveGame/>. Generalised version at <http://demonstrations.wolfram.com/ExpectedDynamicsOfAnImitationModelIn2x2SymmetricGames/>

The reader can replicate the results presented in figure 3 using its [interactive version](#). The differential equation $\dot{x} = x(1-x)(1-2x)$ has three zeros: $x = 0$, $x = 1$ and $x = 0.5$. By calculating the sign of \dot{x} in the domain $[0, 1]$, it is easy to check (see Figure 4) that $x = 0.5$ is an asymptotically stable point whose basin of attraction is the open interval $(0, 1)$. This means that, for large N and any initial state that is neither 0 nor 1, we can expect the state of the system to approach the surroundings of $x = 0.5$ and linger in that area for long.

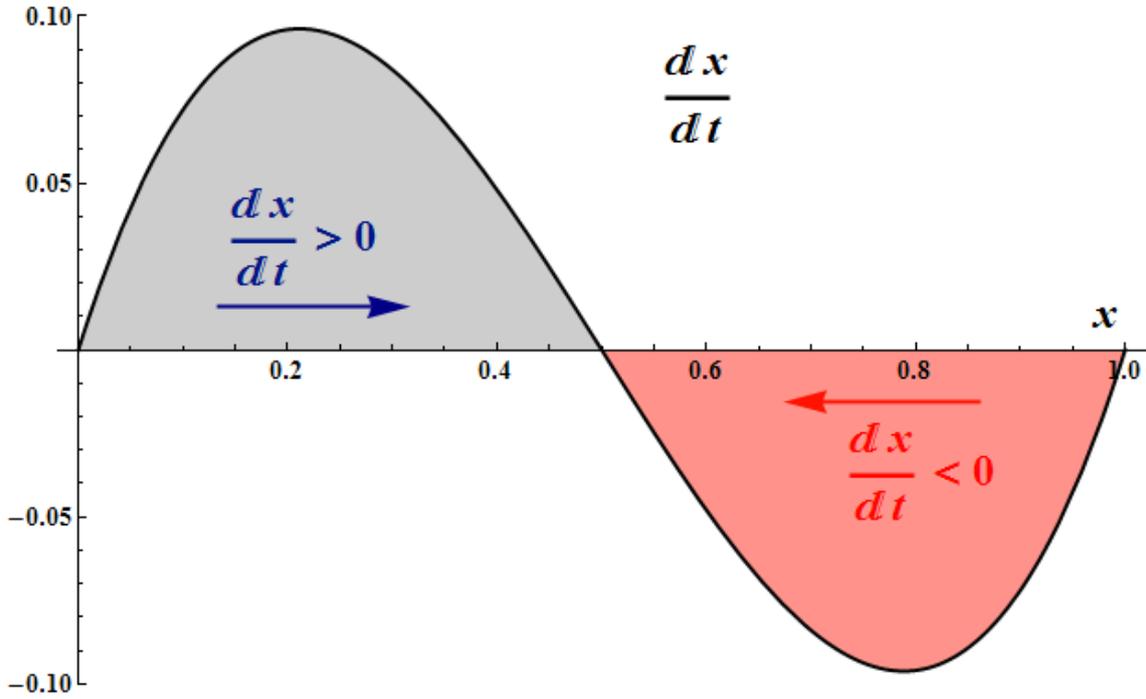


Figure 4. Representation of the function $\dot{x} = x(1-x)(1-2x)$. The value of the time derivative is positive for $0 < x < 0.5$ and negative for $0.5 < x < 1$.

Even though the process will be absorbed in the state $X_n = 1$ with probability one if simulated indefinitely (assuming an initial state $X_0 \neq 0$), the probability of reaching such absorbing state in any given finite time (say, for instance, 3 years running a simulation) can be negligible for a large enough N . This is so even if the simulation time is proportional to the number of agents N : for any given positive constant c , the probability that, after $c \cdot N$ days of computing time, a simulation of the process has reached the absorbing state $X_n = 1$ goes to zero as N grows (assuming an initial state $X_0 \neq 1$). This example illustrates the difference between the long run (a possibly very long but finite time) and ultralong run (the limit as $t \rightarrow \infty$) behaviour of a system [9, 50].

5.2. An intra-population imitation model in a two-population Hawk-Dove game

Let us now suppose that the Hawk-Dove game of example 5.1 is played by individuals belonging to 2 different populations A and B. To keep things simple in this example, let us assume that in each population there are $N (\geq 2)$ individuals and that in time-step n all individuals are randomly matched in couples made up by one individual from population A and one individual from population B.

At the end of every time-step, one randomly selected player from each population revises her strategy, and her revision rule is as follows: “I look at another (randomly selected) individual in my population; if and only if I prefer her last outcome to my last outcome, I adopt her strategy”.

STEP 1: Preliminary simulations

For population size N , consider the step size $\gamma = N^{-1}$. Let $X_{A,n}^\gamma$ be the fraction of individuals in population A (of size N) using strategy H at the beginning of interaction n , and define $X_{B,n}^\gamma$ likewise. Figure 5 shows the evolution of $X_{A,n}^\gamma$ and $X_{B,n}^\gamma$ in two simulations with $N = 10$, for different initial number of H-strategists in each population.

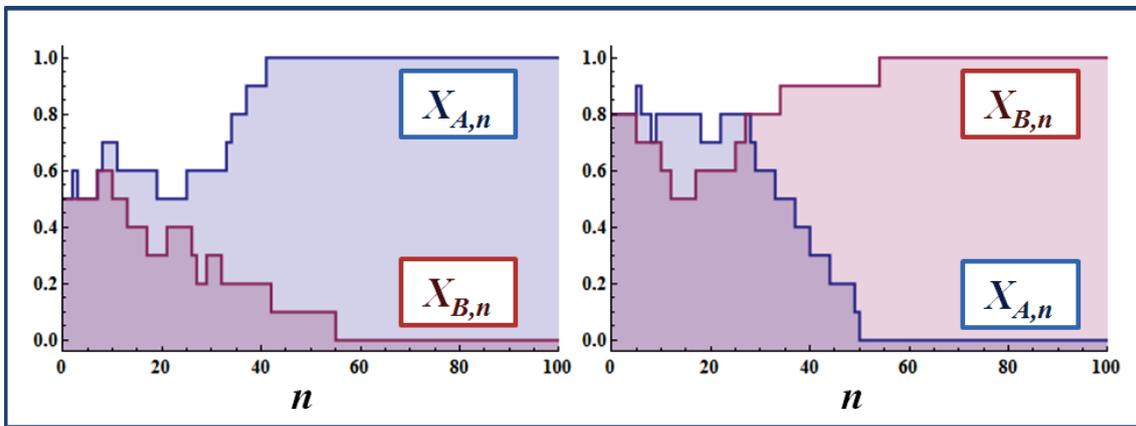


Figure 5. Evolution of the fraction of individuals using strategy H in each population in two simulations with population size = 10 and different initial number of H-strategists in the populations. Interactive figure at

<http://demonstrations.wolfram.com/AnIntraPopulationImitationModelInTheTwoPopulationHawkDoveGame/>. Generalised version at

<http://demonstrations.wolfram.com/AnIntraPopulationImitationModelForInterPopulation2x2Symmetri/>

Most simulations with this model show a quick convergence to a situation in which all individuals in population A play one of the strategies (H or D) and all individuals in population B play the other strategy. The reader can confirm this fact by running simulations with the [interactive version of figure 5](#).

STEP 2 Markov formulation

For populations of size N , this model can be formalised as a Markov process whose state vector in time-step n ($n = 1, 2, \dots$) is $X_n^\gamma = [X_{A,n}^\gamma, X_{B,n}^\gamma]$, where $X_{i,n}^\gamma$ ($i = A, B$) is

the fraction of individuals in population i using strategy H at the beginning of interaction n .

There are $(N + 1)^2$ possible states, depending on the number of individuals using strategy H in each population (this number ranges from 0 to N).

STEP 3 Absorbing states and limiting distribution

In this Markov chain there are precisely 4 absorbing states: $X_n^\gamma = [0, 0]$, $X_n^\gamma = [0, 1]$, $X_n^\gamma = [1, 0]$ and $X_n^\gamma = [1, 1]$, corresponding to the four possible situations in which in each population everyone plays H or everyone plays D. It is straightforward to check that from any non-absorbing state it is possible to reach (at least) one of the absorbing states in a finite number of steps, so the process will necessarily end up in one of the absorbing states. The probability of ending up in each of the different absorbing states depends on the initial conditions, and these probabilities can be estimated by simulation, including confidence intervals [28]. If the state transition matrix is known, the exact probabilities can be calculated analytically. In this example, let $H_{i,n}^N$ be the number of H-strategists in population i ($i = A, B$) in time-step n . Then, for $0 \leq k_A, k_B \leq N$, where k_A and k_B are natural numbers, and given the probabilities of events defined by the revision rule of the model, we can compute the following transition probabilities:

$$P(H_{A,n+1}^N = k_A + 1 | H_{A,n}^N = k_A, H_{B,n}^N = k_B) = \frac{(N - k_A)k_A(N - k_B)}{N^2(N - 1)}$$

$$P(H_{A,n+1}^N = k_A - 1 | H_{A,n}^N = k_A, H_{B,n}^N = k_B) = \frac{k_A k_B (N - k_A)}{N^2(N - 1)}$$

$$P(H_{A,n+1}^N = k_A | H_{A,n}^N = k_A, H_{B,n}^N = k_B) = \frac{k_A^2 + N(N - k_A - 1)}{N(N - 1)}$$

The formulas for $H_{B,n+1}^N$ are analogous. Considering these transition probabilities, for $N = 3$ and numbering the states from 1 to 16 according to the formula: state-number of $[X_{A,n}, X_{B,n}] = 12 \cdot X_{A,n} + 3 \cdot X_{B,n} + 1$, the state transition matrix is

$$P = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 2/3 & 1/3 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 2/3 & 1/3 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 2/3 & 0 & 0 & 0 & 1/3 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1/81 & 2/27 & 2/81 & 0 & 2/27 & 4/9 & 4/27 & 0 & 2/81 & 4/27 & 4/81 & 0 & 0 & 0 & 0 \\ 0 & 2/81 & 4/27 & 4/81 & 0 & 2/27 & 4/9 & 4/27 & 0 & 1/81 & 2/27 & 2/81 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1/3 & 0 & 0 & 0 & 2/3 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 2/3 & 0 & 0 & 0 & 1/3 & 0 & 0 \\ 0 & 0 & 0 & 0 & 2/81 & 2/27 & 1/81 & 0 & 4/27 & 4/9 & 2/27 & 0 & 4/81 & 4/27 & 2/81 & 0 \\ 0 & 0 & 0 & 0 & 0 & 4/81 & 4/27 & 2/81 & 0 & 4/27 & 4/9 & 2/27 & 0 & 2/81 & 2/27 & 1/81 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1/3 & 0 & 0 & 0 & 2/3 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1/3 & 2/3 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1/3 & 2/3 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

The limiting distribution, given an initial state distribution p_0 , is $p_\infty = p_0 \cdot \lim_{n \rightarrow \infty} P^n$, and $\lim_{n \rightarrow \infty} P^n =$

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0.02471 & 0 & 0 & 0.48549 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.48549 & 0 & 0 & 0.00431 \\ 0.00396 & 0 & 0 & 0.84387 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.14821 & 0 & 0 & 0.00396 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0.00396 & 0 & 0 & 0.14821 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.84387 & 0 & 0 & 0.00396 \\ 0.00431 & 0 & 0 & 0.48549 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.48549 & 0 & 0 & 0.02471 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

which provides the probability of ending up in each of the absorbing states for every initial state distribution p_0 (and in particular, for every initial state X_0).

STEP 4 Transient behaviour

If the state transition matrix \mathbf{P} and the initial state X_0 (or the initial state distribution) are known, it is easy to calculate the state distribution corresponding to any time step.

However, the dimension of the state transition matrix \mathbf{P} is $(N+1)^2 \times (N+1)^2$, so even for moderate values of N it can be unfeasible to work with the matrix \mathbf{P} . For large values of N we can expect stochastic approximation theory to be useful. In our current case $X_n^\gamma \in I = [0, 1] \times [0, 1]$ and there is the relation:

$$\Delta X_n^\gamma = \frac{1}{N} \Delta H_n^N = \gamma \cdot Y_n^\gamma$$

where $\gamma = N^{-1}$ is a scalar number that is small for large N , and $Y_n^\gamma = \Delta H_n^N$ is a random vector whose components $(\Delta H_{A,n}^N, \Delta H_{B,n}^N)$ may take the values -1, 1 or 0 with a probability distribution (see step 3) that depends only on the current state X_n^γ and on the population size N (or, equivalently, on the step size $\gamma = N^{-1}$). Let us now check that this model fulfils the conditions outlined in section 4:

1.- For every positive natural number N , the state space of the Markov process X_n^γ belongs to the closed convex set $I = [0, 1] \times [0, 1]$, and the minimum distance from any fixed point in I to some point in the state space of X_n^γ goes to zero as γ goes to 0 (i.e., as N grows).

2.- The expected value of Y_n^γ conditioned on $X_n^\gamma = x = [x_A, x_B]$ is (omitting the superscripts in the variables):

$$\begin{aligned} g(x, \gamma) &= E(Y_n | X_n = [x_A, x_B]) = E(\Delta H_n | X_n = [x_A, x_B]) = \\ &= \frac{N}{N-1} [x_A(1-x_A)(1-2x_B), x_B(1-x_B)(1-2x_A)] = \\ &= (1 + \frac{1}{N-1}) [x_A(1-x_A)(1-2x_B), x_B(1-x_B)(1-2x_A)] = \\ &= [x_A(1-x_A)(1-2x_B), x_B(1-x_B)(1-2x_A)] + O(\frac{1}{N}) = \\ &= [x_A(1-x_A)(1-2x_B), x_B(1-x_B)(1-2x_A)] + O(\gamma) \end{aligned}$$

uniformly in x .

The function $g(x_A, x_B) = [x_A(1 - x_A)(1 - 2x_B), x_B(1 - x_B)(1 - 2x_A)]$ is made up of polynomials, and therefore it is differentiable on the compact set I and its derivative matrix $g'(x)$ is also made up of polynomials, so it is bounded and Lipschitz on I .

3.- $s(x, \gamma) = E[(Y_n^\gamma - g(x, \gamma))^{Tr} \cdot (Y_n^\gamma - g(x, \gamma)) | X_n^\gamma = x]$ is made up of sums of products of polynomials, each of which uniformly converges to a polynomial in x . Consequently, $s(x)$ is Lipschitz on I .

4.- $E(|Y_n^\gamma|^3 | X_n^\gamma = x)$ is uniformly bounded, because each component of $|Y_n^\gamma| = |\Delta H_n^N|$ is bounded by 1.

Therefore the results described in section 4 apply.

Consider the associated limit mean ODE for X_n^γ : the differential equation

$$\dot{x} = [x_A(1 - x_A)(1 - 2x_B), x_B(1 - x_B)(1 - 2x_A)]$$

with $0 \leq x_A, x_B \leq 1$.

Figure 6 shows several trajectories for this differential equation. For large N (i.e. small γ), we can expect the state of the process X_n^γ to evolve closely to the trajectory of the ODE that goes through the initial state x_0 . We can also expect the state of the process in time-step $n = \text{int}(t \cdot N)$ to get close to the solution at time t of that same ODE. The reader can confirm this fact by running simulations with the [interactive version of figure 6](#).

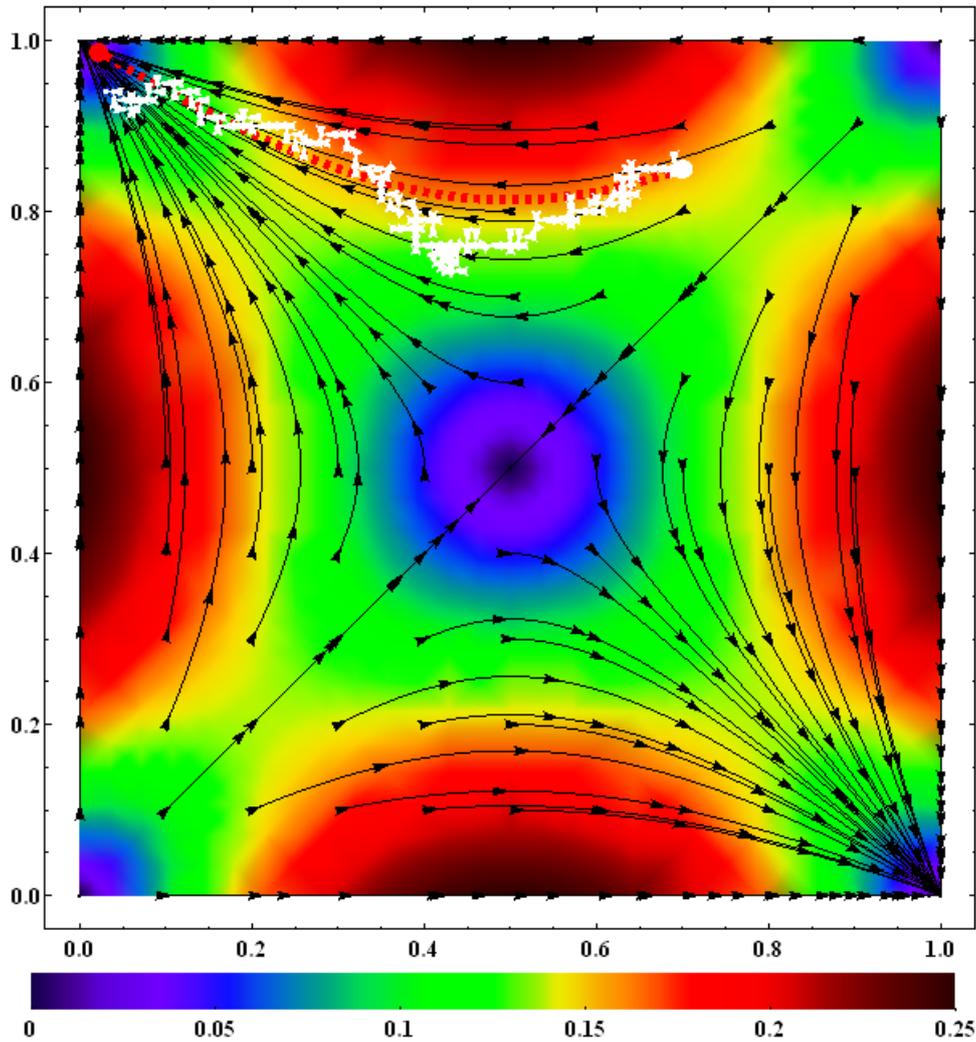


Figure 6. This figure shows several trajectories of the differential equation that describes the expected move of $X_n = [X_{A,n}, X_{B,n}]$ for large values of N (the mean field). In dashed red, the figure shows the trajectory $x(t)$ with initial conditions $x(t=0) = [0.7, 0.85]$, from $t=0$ to $t=6$. In white, the figure shows the results of a simulation of the stochastic process with $N=100$ and initial conditions $X_0 = [0.7, 0.85]$, from $n=0$ to $n=t \cdot N=600$. Interactive figure at <http://demonstrations.wolfram.com/ExpectedDynamicsOfAnIntraPopulationImitationModelInTheTwoPop/>. Generalised version at <http://demonstrations.wolfram.com/ExpectedDynamicsOfAnIntraPopulationImitationModelForInterPop/>

5.3. Generalizations for arbitrary payoffs

All the results in sections 5.1 and 5.2 about the transient behaviour of the imitation models in the Hawk-Dove game can be generalised to any 2x2 game. Consider, for instance, the symmetric case, with the notation shown below for the preferences over outcomes (payoffs) for the row player, with higher payoff denoting higher preference:

	H	D
H	hh	hd
D	dh	dd

Following the same reasoning as in example 5.1, the mean-field equation with arbitrary payoffs for the one-population model described in section 5.1 reads:

$$g(x) = x(1-x) [\delta_1 x^2 + \delta_2 (1-x)^2 + \delta_3 (1-x)x]$$

where x is the proportion of H-players in the population,

$$\delta_1 = \text{Sign}(hh - dh),$$

$$\delta_2 = \text{Sign}(hd - dd), \text{ and}$$

$$\delta_3 = \text{Sign}(hd - dh) + \text{Sign}(hh - dd).$$

Some examples of these dynamics can be explored using the generalized versions of [interactive figure 2](#) and [interactive figure 3](#).

Similarly, the mean-field equation with arbitrary payoffs for the two-population model described in section 5.2 reads:

$$g(x_A, x_B) = [g_1(x_A, x_B), g_2(x_A, x_B)]$$

$$g_1(x_A, x_B) = x_A(1-x_A) [\delta_1 x_B^2 + \delta_2 (1-x_B)^2 + \delta_3 (1-x_B)x_B]$$

$$g_2(x_A, x_B) = x_B(1-x_B) [\delta_1 x_A^2 + \delta_2 (1-x_A)^2 + \delta_3 (1-x_A)x_A]$$

where x_A is the proportion of H-players in population A and x_B is the proportion of H-players in population B. Some examples of these dynamics can be explored using the generalized versions of [interactive figure 5](#) and [interactive figure 6](#).

In the following sub-section we prove convergence results to the mean dynamic for an even broader family of models, namely the intra-population imitation model in n -player games with arbitrary payoffs. In particular, this family encompasses all the models presented above.

5.4. The intra-population imitation model in n -player games

The results about convergence to the mean dynamic for the imitation stochastic processes described above can be extended to any n -player game with a finite number of strategies, either in one single population, or considering that each player's position in the game corresponds to one individual selected from a different population.

Assuming that the size of each population is N_i , with N_i proportional to the size N of the smallest population, and that, after each time-step, k random players (or k_i players in each population, with $k_i \leq k$) are given the opportunity to revise their strategy according

to the previous imitation rule, the process for a given N can be characterised as a Markov chain with a state vector X_n^γ that contains the proportions of each type of pure-strategy-player in each population.

The process is such that $\Delta X_n^\gamma = \gamma \cdot Y_n^\gamma$, where $\gamma = N^{-1}$ is a scalar number that is small for large N , and Y_n^γ is a random vector whose components may take bounded integer values (between $-k$ and $+k$) with a probability distribution that depends only on the current state $X_n^\gamma = x$ and on the population size N . Note that:

1.- For every positive natural number N , the state space of the Markov process X_n^γ belongs to the closed convex polyhedron I : each component of the state vector belongs to $[0, 1]$, and the components that correspond to the strategies available in each population add up to one. The minimum distance from any fixed point in I to some point in the state space of X_n^γ goes to zero as γ goes to 0 (i.e., as N grows).

Considering that the probability that a player changes from one strategy to another is a function of the form $\text{Poly}(x) + O(\frac{1}{N})$ uniformly in x , and that, for $x \in I$, the sum of products of functions of this type is another function of this kind, we can state the following:

2.- The expected value of Y_n^γ conditioned on $X_n^\gamma = x$ is $g(x, \gamma) = g(x) + O(\frac{1}{N})$ uniformly in x , where $g(x)$ is made up by polynomials, and therefore it is differentiable on the compact set I and its derivative matrix $g'(x)$ is also made up by polynomials, so it is bounded and Lipschitz on I .

3.- $s(x, \gamma) = E[(Y_n^\gamma - g(x, \gamma))^{Tr} \cdot (Y_n^\gamma - g(x, \gamma)) | X_n^\gamma = x]$ is made up by sums of products of polynomials, each of which uniformly converges to a polynomial in x . Consequently, $s(x)$ is Lipschitz on I .

4.- $E(|Y_n^\gamma|^3 | X_n^\gamma = x)$ is uniformly bounded, because each component of $|Y_n^\gamma|$ is bounded by k .

Therefore the results described in section 4 apply.

6 Conclusions

This paper has discussed the relation between the actual behaviour of computer simulations of some simple *stochastic* models and their expected *deterministic* behaviour where the inherent stochasticity of the model is replaced by a mean dynamic. We have presented selected results from stochastic approximation theory emphasising those aspects that can be particularly useful to understand the dynamics of computer simulations. The framework presented here can be useful to analyse relatively simple models with a large number of interacting individuals or where individuals modify their state by small steps, as in most of the referenced works. The potential of this technique has been illustrated through the analysis of two particular cases, which neatly show how the qualitative dynamics of simulated models of interacting elements can vary drastically with the number of elements in the system. Finally, we have provided convergence results for simulations of a broad family of models of evolutionary games.

Acknowledgements

The authors gratefully acknowledge financial support from the Spanish Ministry of Education (JC2009-00263) and MICINN (CONSOLIDER-INGENIO 2010: CSD2010-00034, and DPI2010-16920). We also thank two anonymous reviewers for their useful comments.

Appendix. Approximation of difference equations by differential equations

This appendix discusses the relation between a discrete-time difference equation of the form $\Delta x_n = \gamma f(x_n)$, with initial point x_0 , and the solution $x(t, x_0)$ of its corresponding continuous time differential equation $\dot{x} = f(x)$ with $x(t = 0) = x_0$.

This relationship can be neatly formalised using Euler's method [12]. If $f(x)$ satisfies some conditions that guarantee convergence [23, 35], e.g. if $f(x)$ is globally Lipschitz or if it is smooth on a forward-invariant compact domain of interest D , the difference between the solution $x(t, x_0)$ of the differential equation at time $t = T < \infty$ and the solution x_n of the equation in differences at step $n = \text{int}(T/\gamma)$ converges to 0 as the step-size parameter γ tends to 0:

$$x_n = x_0 + \sum_{i=0}^{\text{int}(\frac{T}{\gamma})} \gamma \cdot f(x_i) \xrightarrow{\gamma \rightarrow 0} x(T, x_0)$$

and, for every T , $\max_{n=0,1,\dots,N} \|x_n - x(n \frac{T}{N}, x_0)\| \xrightarrow{N \rightarrow \infty} 0$

As an example, consider a vector $x = [x_1, x_2]$, the function $f(x) = [x_2, -x_1]$, the differential equation $\dot{x} = f(x)$ and its associated (deterministic) equation in differences $\Delta x_n = \gamma f(x_n)$. Figure 7 shows a trajectory map of the differential equation $\dot{x} = f(x)$, together with several values of discrete processes that follow the equation in differences $\Delta x_n = \gamma f(x_n)$, for different values of γ and for the same initial value $x_0 = [x_1, x_2]_0$. It can be seen how, for decreasing values of γ and for a correspondingly increasing finite number of steps $n = \text{int}(T/\gamma)$, the discrete process gets closer and closer to the trajectory of the differential equation that goes through x_0 . The reader can confirm this fact by running simulations with the [interactive version of figure 7](#).

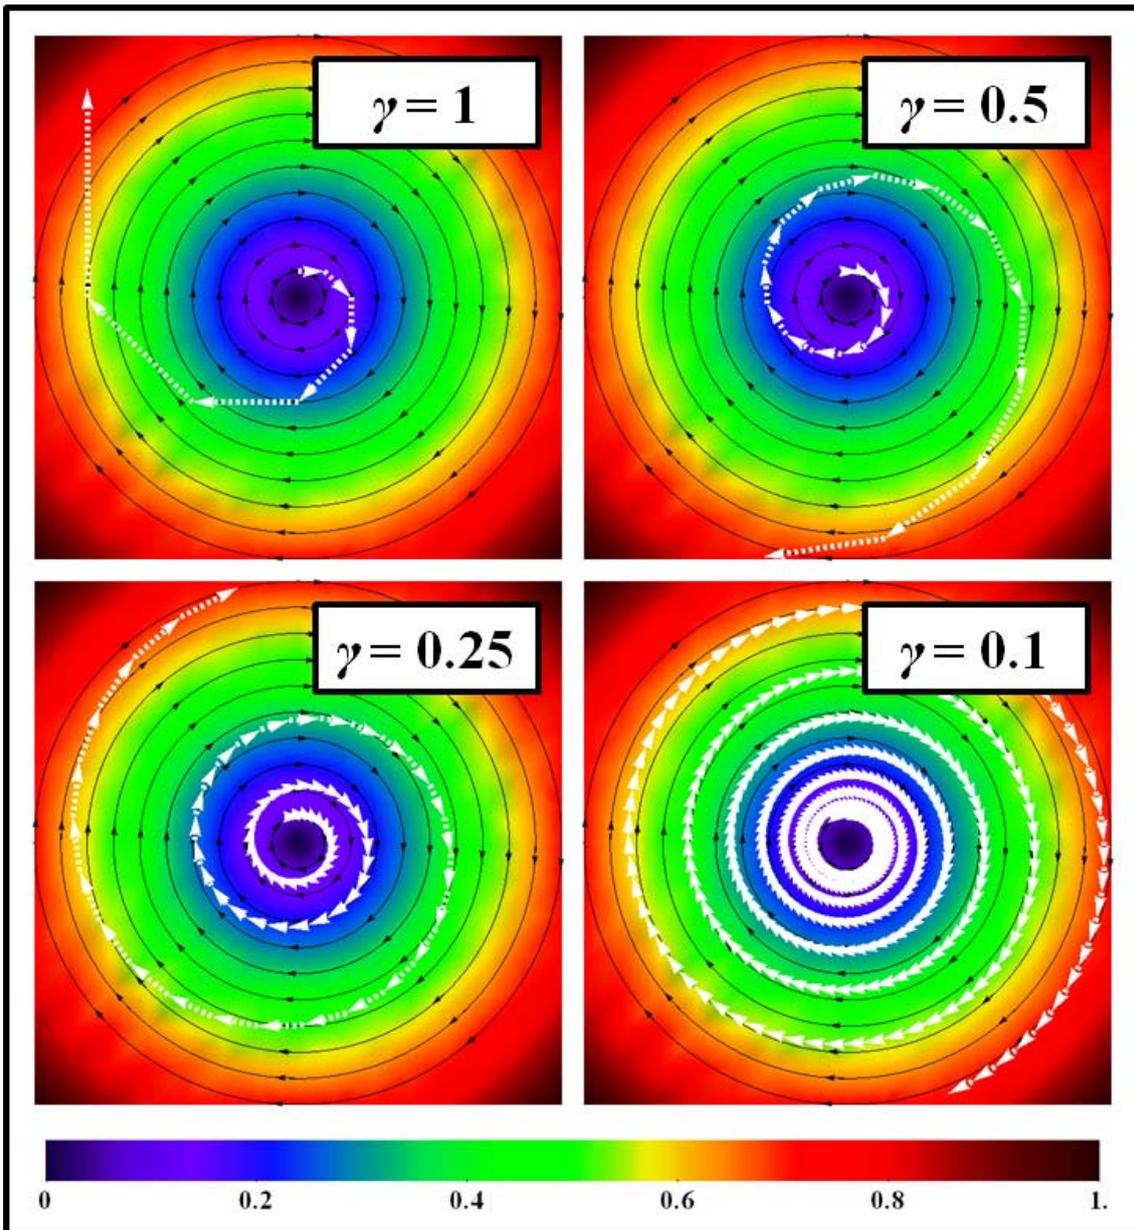


Figure 7. Convergence of difference and differential equations for small step-size. Let $x = [x_1, x_2]$ be a generic point in the real plane. Figure 7 shows a trajectory map of the differential equation $\dot{x} = f(x) = [x_2, -x_1]$, together with several values of the discrete process that follow the equation in differences $x_{n+1} - x_n = \gamma f(x_n)$, for different values of the step-size parameter γ and for a chosen initial value $x_0 = [x_1, x_2]_0$. The background is coloured using the norm of the expected motion, rescaled to be in the interval $(0, 1)$. It can be seen how, for decreasing values of γ , the discrete process tends to follow temporarily the trajectory of the differential equation that goes through x_0 . Interactive figure at <http://demonstrations.wolfram.com/DifferenceEquationVersusDifferentialEquation/>.

References

- [1] Banisch, S., Lima, R., Araújo, T.: Agent based models and opinion dynamics as Markov chains. *Social Networks* (in Press)
- [2] Barrat, A., Barthélemy, M., Vespignani, A.: *Dynamical Processes on Complex Networks*. Cambridge University Press (2008)
- [3] Beggs, A.W.: Large deviations and equilibrium selection in large populations. *J. Econ. Theory* **132**(1), 383-410 (2007)
- [4] Beggs, A.W.: On the convergence of reinforcement learning. *J. Econ. Theory* **122**(1), 1-36 (2005)
- [5] Beggs, A.W.: Stochastic evolution with slow learning. *J. Econ. Theory* **19**(2), 379-405 (2002)
- [6] Benaïm, M., Le Boudec, J.-.: A class of mean field interaction models for computer and communication systems. *Perform Evaluation* **65**(11-12), 823-838 (2008)
- [7] Benaim, M., Weibull, J.W.: Deterministic approximation of stochastic evolution in games. *Econometrica* **71**(3), 873-903 (2003)
- [8] Benveniste, A., Priouret, P., Metivier, M.: *Adaptive algorithms and stochastic approximations*. Springer-Verlag New York, Inc (1990)
- [9] Binmore, K.G., Samuelson, L., Vaughan, R.: Musical Chairs: Modeling Noisy Evolution. *Games Econ. Behav.* **11**(1), 1-35 (1995)
- [10] Börgers, T., Sarin, R.: Learning through reinforcement and replicator dynamics. *J. Econ. Theory* **77**(1), 1-14 (1997)
- [11] Borkar, V.S.: *Stochastic approximation: a dynamical systems viewpoint*. Cambridge University Press, New York (2008)
- [12] Borrelli, R.L., Coleman, C.S.: *Differential equations: a modeling approach*. Prentice-Hall, Englewoods Cliffs (1987)
- [13] Boylan, R.T.: Laws of large numbers for dynamical systems with randomly matched individuals. *J. Econ. Theory* **57**(2), 473-504 (1992)
- [14] Boylan, R.T.: Continuous Approximation of Dynamical Systems with Randomly Matched Individuals. *J. Econ. Theory* **66**(2), 615-625 (1995)
- [15] Erev, I., Roth, A.E.: Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria. *Am. Econ. Rev.* **88**(4), 848-881 (1998)
- [16] Ethier, S.N., Kurtz, T.G.: *Markov Processes: Characterization and Convergence*. Wiley Series in Probability and Statistics (2005)
- [17] Fudenberg, D., Kreps, D.M.: Learning Mixed Equilibria. *Game. Econ. Behav.* **5**(3), 320-367 (1993)
- [18] Fudenberg, D., Levine, D.K.: *The theory of learning in games*. The MIT Press (1998)
- [19] Galán, J.M., Latek, M.M., Rizi, S.M.M.: Axelrod's Metanorm Games on Networks. *PLoS ONE* **6**(5), e20474 (2011)
- [20] Gilbert, N., Troitzsch, K.G.: *Simulation for the social scientist*. McGraw-Hill (2005)
- [21] Gintis, H.: Agent-Based Models as Markov Processes: Theory and Application to the Dynamics of Market Exchange. *ACM TIST* (in Press)
- [22] Gleeson, J.P., Melnik, S., Ward, J.A., Porter, M.A., Mucha, P.J.: Accuracy of mean-field theory for dynamics on real-world networks. *Phys Rev E.* **85**(2) (2012)

- [23] Hairer, E., Nørsett, S., Wanner, G.: Solving Ordinary Differential Equations I: Nonstiff Problems. Springer (1993)
- [24] Hofbauer, J., Sandholm, W.H.: On the global convergence of stochastic fictitious play. *Econometrica* **70**(6), 2265-2294 (2002)
- [25] Hopkins, E.: Two competing models of how people learn in games. *Econometrica* **70**(6), 2141-2166 (2002)
- [26] Huet, S., Deffuant, G.: Differential Equation Models Derived from an Individual-Based Model Can Help to Understand Emergent Effects. *J. Artif. Soc. Soc. Simulat.* **11**(2), 10 (2008)
- [27] Izquierdo, L.R., Izquierdo, S.S., Gotts, N.M., Polhill, J.G.: Transient and asymptotic dynamics of reinforcement learning in games. *Game. Econ. Behav.* **61**(2), 259-276 (2007)
- [28] Izquierdo, L., Izquierdo, S., Galán, J.M., Santos, J.I.: Techniques to Understand Computer Simulations: Markov Chain Analysis. *J. Artif. Soc. Soc. Simulat.* **12**(1), 6 (2009)
- [29] Izquierdo, S.S., Izquierdo, L.R., Gotts, N.M.: Reinforcement learning dynamics in social dilemmas. *J. Artif. Soc. Soc. Simulat.* **11**(2), 1 (2008)
- [30] Izquierdo, S.S., Izquierdo, L.R., Vega-Redondo, F.: The option to leave: Conditional dissociation in the evolution of cooperation. *J. Theor. Biol.* **267**(1), 76-84 (2010)
- [31] Kotliar, G., Savrasov, S.Y., Haule, K., Oudovenko, V.S., Parcollet, O., Marianetti, C.A.: Electronic structure calculations with dynamical mean-field theory. *Rev Mod Phys* **78**(3), 865-951 (2006)
- [32] Kulkarni, V.G.: Modeling and Analysis of Stochastic Systems. Chapman & Hall/CRC (1995)
- [33] Kushner, H.: Stochastic approximation: A survey. *WIREs Computation Stat* **2**(1), 87-96 (2010)
- [34] Kushner, H.J., Yin, G.G.: Stochastic Approximation Algorithms and Applications. Springer-Verlag. New York (2003)
- [35] Lambert, M.F.: Numerical Methods for Ordinary Differential Systems. John Wiley & Sons Ltd, Chichester (1991)
- [36] Ljung, L.: Analysis of Recursive Stochastic Algorithms. *IEEE T Automat Contr* **AC-22**(4), 551-575 (1977)
- [37] López-Pintado, D.: Diffusion in complex social networks. *Game. Econ. Behav.* **62**(2), 573-590 (2008)
- [38] Macy, M.W., Flache, A.: Learning dynamics in social dilemmas. *Proc. Natl. Acad. Sci. U. S. A.* **99**(3), 7229-7236 (2002)
- [39] Morozov, A., Poggiale, J.C.: From spatially explicit ecological models to mean-field dynamics: The state of the art and perspectives. *Ecol. Complex.* **10**, 1-11 (2012)
- [40] Norman, M.F.: Markov Processes and Learning Models. Academic Press, New York (1972)
- [41] Norman, M.F.: Some convergence theorems for stochastic learning models with distance diminishing operators. *J. Math. Psychol.* **5**(1), 61-101 (1968)
- [42] Perc, M., Szolnoki, A.: Coevolutionary games-A mini review. *BioSystems* **99**(2), 109-125 (2010)
- [43] Rozonoer, L.I.: On deterministic approximation of Markov processes by ordinary differential equations. *Math. Probl. Eng.* **4**(2), 99-114 (1998)
- [44] Sandholm, W.H.: Stochastic imitative game dynamics with committed agents. *J. Econ. Theory* **147**(5), 2056-2071 (2012)

- [45] Sandholm, W.H.: Deterministic Evolutionary Dynamics. In: Durlauf, S. N. and Blume, L. E. (eds.) The New Palgrave Dictionary of Economics v.2. Palgrave Macmillan (2008)
- [46] Szabó, G., Fáth, G.: Evolutionary games on graphs. Phys Rep **446**(4-6), 97-216 (2007)
- [47] Vega-Redondo, F.: Complex Social Networks. Cambridge University Press, Cambridge (2007)
- [48] Vega-Redondo, F.: Economics and the Theory of Games. Cambridge University Press, Cambridge (2003)
- [49] Weiss, P.: L'hypothèse du champ moléculaire et la propriété ferromagnétique. J. de Phys. **6**, 661-690 (1907)
- [50] Young, H.P.: Individual Strategy and Social Structure: An Evolutionary Theory of Institutions. Princeton University Press (1998)