Evolutionary Dynamics of Resource Allocation in the Colonel Blotto Game

Damián G. Hernández & Damián H. Zanette

Journal of Statistical Physics
1

ISSN 0022-4715
Volume 151
Combined 3-4

J Stat Phys (2013) 151:623-636
DOI 10.1007/s10955-012-0659-7

Springer

J Stat Phys (2013) 151:623–636
DOI 10.1007/s10955-012-0659-7

Evolutionary Dynamics of Resource Allocation in the Colonel Blotto Game

Damián G. Hernández · Damián H. Zanette

Received: 23 July 2012 / Accepted: 28 November 2012 / Published online: 14 December 2012

Abstract We provide an evolutionary game-theoretical formulation for a model of resource allocation—the Colonel Blotto game. In this game, two players with different total resources must entirely distribute them among a set of items. Each item is won by the player that assigned higher resources to it, and the payoff of each player is the total number of won items. Our evolutionary formulation makes it possible to obtain optimal strategies as the equilibrium states of a dynamical process. At the same time, it naturally requires considering a population of players—whose strategies evolve by imitation and random fluctuations—thus better approaching a realistic situation with many economic agents. Results show, in particular, how agents with low total resources manage to maximize their winnings in spite of their intrinsically disadvantageous condition.

Keywords Socio-economic dynamics · Evolutionary games · Resource management

1 Introduction

Resource allocation is a basic ingredient in the management of a broad class of economic systems, whereby resources are distributed among different components and/or stages of a process in order to—hopefully—maximize its yield. The complexity of resource allocation recognizes at least two origins, involving competition at different levels. Internally to the process, it is necessary to decide how to divide and where to assign the available resources to improve efficiency as much as possible. Externally, such decision may be modulated by the existence of other economic actors engaged in similar activities, thus competing with each other. For instance, with the aim of increasing profits, a company elaborating an edible product—hamburgers, say—must make a decision as to whether raise or lower the quality

D.G. Hernández · D.H. Zanette (✉)
Centro Atómico Bariloche, Consejo Nacional de Investigaciones Científicas y Técnicas,
8400 San Carlos de Bariloche, Río Negro, Argentina
e-mail: zanette@cab.cnea.gov.ar

D.G. Hernández
e-mail: damian.g.h.1@gmail.com

of meat, or to finance better publicity, or to improve its distribution network, and this decision must necessarily take into account the corresponding strategies of other hamburger producers.

However, the problem of resource allocation is not limited to the purely economic domain. Just to give a familiar example, the number of hours per week a scientist decides to devote to each one of his/her collaboration projects—taking into account the finiteness of the available time, the intrinsic difficulty of each project, and its foreseen impact in the community—is a problem of exactly the same kind. Within a different realm, biological evolution optimizes reproductive success by allocating the resources of an organism, in intraspecific competition with its likes, either increasing the probability of mating by a more efficient courtship, or making the number of offsprings larger, or enhancing dedication to parental care, among many other strategies.

Due to the role of competition in the process of resource allocation, game theory provides a natural frame to give the problem a mathematical formulation. A classical, well-studied setting is the Colonel Blotto game. In this auction-like game, two players must distribute their total resources among a fixed number of items. Each item is won by the player who allocated the higher amount for it, and the player's total payoff is the number of won items. An interesting aspect of this game is that it allows considering non-identical players, i.e. players with different total resources, and thus analyzing how optimal strategies are adapted to the disparate situation of each opponent. The fancy name of Colonel Blotto's comes from a setup of the game where two armies must distribute their forces among a certain number of battlefields.

Apparently, the first version of the Colonel Blotto game was formulated by the famous French mathematician Émile Borel in 1921 [1], and optimal strategies for the case of three items and identical players were found by Borel and Ville several years later [2]. In 1950, Gross and Wagner provided a solution for an arbitrary number of items and identical players, and a solution for two items and non-identical players [3]. Much more recently, variants of the game were solved in more general situations [4], and applications to real-life problems such as electoral campaigns [5–7] appeared. Roberson found the first complete solution to the original Colonel Blotto game for arbitrary number of items and non-identical players—which we partly review in our Sect. 2.1—just a few years ago [8]. Golman and Page discussed the possibility of attaining mixed strategies through dynamic learning, numerically implemented using the replicator dynamics [9]. These, and a few further contributions [10, 11], were recently summarized and reviewed [12].

In this paper, we put the Colonel Blotto game in the framework of evolutionary game theory. In contrast with traditional game theory, its evolutionary version regards optimized strategies as the result of a dynamical process where players tend to maximize the individual winnings by modifying their own strategies as the process goes on [13, 14]. This scenario is inspired in biological (Darwinian) evolution, where strategies vary from parents to their descendence by genetic mutations under the pressure of natural selection. Evolutionary game theory, thus, has the advantage of highlighting the role of change with adaptation, opposite in a sense to the static view of strategy optimization adopted by the traditional theory. Additionally, evolutionary game theory naturally leads to considering a whole population of players engaged in the same game—with a multitude of agents able to adopt different strategies, which compete with each other—and is therefore better adapted to provide realistic models of social phenomena.

In the next section, we first review previous results for optimal strategies in the Colonel Blotto game with two players, as provided by traditional game theory. We then generalize these results to a population of players of two types, with low and high total resources, where each player can confront an opponent of either type. In Sect. 3 we introduce the evolutionary model and present numerical results which validate the assumptions of our previous analytical study. Numerical simulations of the evolution process make as well possible to determine what strategies are actually adopted for parameter sets where more than one analytical solution coexist. Finally, we discuss a few generalizations—which emphasize the robustness of the main conclusions—and summarize our contribution.

2 Optimal Allocation Strategies in the Colonel Blotto Game

2.1 The Two-Player Game

In the traditional formulation of the Colonel Blotto game [8], two players A and B have total resources X_A and X_B, respectively. Without generality loss, we take $X_A \le X_B$. Both players must distribute their whole resources among m items. The items are equivalent to each other, in the sense that the players do not have *a priori* preferences as to which item should be assigned more resources. Let x_A^j and x_B^j be the resources allocated by each player to item j $(j = 1, \ldots, m)$. The player's payoffs, π_A and π_B, are respectively given by the fraction of items for which $x_A^j > x_B^j$ and the fraction of items for which $x_B^j > x_A^j$, namely,

$$\pi_{A,B} = \frac{1}{m} \big\| \{ j / x_{A,B}^j > x_{B,A}^j \} \big\|, \tag{1}$$

where $\| \cdot \|$ indicates cardinality. As $x_{A,B}^j$ varies continuously over the interval $[0, X_{A,B}]$ for all j, the marginal cases with $x_A^j = x_B^j$ can be disregarded. Since $\pi_A + \pi_B = 1$ for any resource distribution, this is a constant-sum game. For this kind of game, optimal strategies coincide with Nash equilibria, where each player's payoff is maximized given the strategy of the opponent.

If $X_B > m X_A$, the game has a set of trivial optimal strategies, with $x_B^j = X_B/m$ for all j. In this situation, player B always wins all items ($\pi_B = 1$). On the other hand, when $X_B < m X_A$ Nash equilibria do not correspond to pure strategies, i.e. to strategies where a fixed resource amount is allocated to each item [8]. Under this condition, in fact, any pure strategy is outdone by another pure strategy where resources are taken from a single item and distributed among the others. In this case, optimal strategies are mixed, and probabilities are assigned to each possible resource distribution for each player.

Let $P_i(x_i^1, \ldots, x_i^m)$ be the probability distribution for the set of resources $\{x_i^1, \ldots, x_i^m\}$ allocated by player $i \in \{A, B\}$. We denote by $P_i^j(x_i^j)$ the marginal probability distribution for the resources allocated to item j,

$$P_i^j\big(x_i^j\big) = \int dx_i^1 \cdots \int dx_i^{j-1} \int dx_i^{j+1} \cdots \int dx_i^m \, P_i\big(x_i^1, \ldots, x_i^m\big), \tag{2}$$

and $F_i^j(x_i^j)$ the corresponding cumulated distribution,

$$F_i^j\big(x_i^j\big) = \int_0^{x_i^j} dx \, P_i^j(x). \tag{3}$$

The expected contribution to each player's payoff from item j, $\pi_{A,B}^j$, is the probability that the corresponding allocated resources are larger than those of the opponent, i.e.

$$\pi_{A,B}^j = \int_0^{X_{A,B}} dx \, P_{A,B}^j(x) \int_0^x dx' \, P_{B,A}^j\big(x'\big) \equiv \int_0^1 F_{B,A}^j \, dF_{A,B}^j. \tag{4}$$

The expected total payoff is $\pi_{A,B} = m^{-1} \sum_j \pi_{A,B}^j$.

The maximization of $\pi_{A,B}$ with respect to the marginal distributions $P^j_{A,B}(x)$, which determines the Nash equilibria, must be performed with the condition that the sum of the allocated resources equals the total resources $X_{A,B}$. In terms of the distributions, this condition reads

$$X_{A,B} = \sum_j \int_0^{X_{A,B}} dx\, x P^j_{A,B}(x) \equiv \sum_j \int_0^1 x\, dF^j_{A,B}. \tag{5}$$

Introducing Lagrange multipliers $\lambda_{A,B}$, it is possible to write the two Lagrangians [4, 8]

$$\mathcal{L}_{A,B} = \frac{1}{m} \sum_j \int_0^1 F^j_{B,A} dF^j_{A,B} + \lambda_{A,B} \left(X_{A,B} - \sum_j \int_0^1 x\, dF^j_{A,B} \right), \tag{6}$$

whose joint extremization with respect to the set of cumulative distributions F^j_A and F^j_B provides the optimal strategies, as solutions of the corresponding Euler–Lagrange equations.

For $m \gg 1$—a limit to which we stick from now on—the result for the marginal probability distributions can be expressed in a particularly compact form by introducing the rescaled variables $u^j_{A,B} = (2X_B/m)^{-1} x^j_{A,B}$ and the total resource ratio $\alpha = X_A/X_B$ (≤ 1). It reads [8]

$$P_A(u_A) = (1-\alpha)\delta(u_A) + \alpha U(u_A; 0, 1) \tag{7}$$

and

$$P_B(u_B) = U(u_B; 0, 1). \tag{8}$$

Because of the *a priori* equivalence of all items, these distributions are the same for all $j = 1, \dots, m$. For the sake of clarity, thus, we have dropped the index j. In the above expressions, $\delta(u)$ is Dirac's delta and $U(u; u_1, u_2) = 1$ for $u \in (u_1, u_2)$ and 0 otherwise.

In terms of the original variables $x_{A,B}$, for non-zero allocated resources—specifically, for $x_A \neq 0$—the two distributions are flat, and differ from zero up to a cutoff at $x_{A,B} = 2X_B/m$. In the optimal strategy, therefore, every positive amount of allocated resources, up to a maximum of $2X_B/m$, is equally probable for both players. The delta-like first term in the right-hand side of Eq. (7), however, represents a finite probability that player A allocates no resources to one or more items. This is due to the relative disadvantage of A, whose total resources are less than those of B, and is thus forced to resign winning some of the items in order to get a maximal payoff. The average payoffs of the two players are, respectively, $\langle \pi_A \rangle = \alpha/2$ $(\leq 1/2)$ and $\langle \pi_B \rangle = 1 - \alpha/2$ $(\geq 1/2)$. As expected, $P_A \equiv P_B$ for $\alpha = 1$, i.e. when $X_A = X_B$ and the two players become indistinguishable from each other.

2.2 A Population Playing the Game

Our aim is now to extend the above results for two players to the case of a population formed by players of two types, A and B. The total resources per player are X_A and X_B, respectively, with $X_A \leq X_B$. Each member of the population is allowed to play the game against an opponent of any type. There are, therefore, three possible kinds of matches: A vs. A, A vs. B, and B vs. B. For a given player, we denote by ϵ and $1-\epsilon$ the probabilities that the opponent is, respectively, of the same type and of different type. The game rules are the same as before, namely, the total resources of both players are distributed among all items, and each item is won by the player who allocated the largest amount.

To find the optimal strategies in this case, it would be in principle necessary to consider that each player in the population may have a different strategy, described by its own

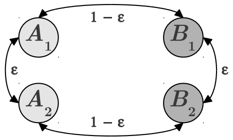

Fig. 1 Schematic four-player representation of the whole population. Each player, of type A or B, confronts a player of the same type with probability ϵ and a player of different type with probability $1-\epsilon$

marginal probability distribution. This would lead to the formidable problem of simultaneously optimizing as many strategies as the population size. Instead of dealing with this, we make in the following a drastic simplification of the problem. Its validity is later verified by means of numerical simulations, as explained in Sect. 3.

We represent the whole population as a set of just four players, A_1, A_2, B_1 and B_2, the former two belonging to type A and the latter two to type B. Figure 1 schematizes their possible interactions, with the corresponding probabilities. Assigning cumulative probability distributions $F^j_{A_1}$, $F^j_{A_2}$, $F^j_{B_1}$, and $F^j_{B_2}$ to the resources allocated by each player to item j, it is straightforward to realize that the two Lagrangians in Eq. (6) are now replaced by four Lagrangians, introducing four Lagrange multipliers:

$$\mathcal{L}_{A_1,B_1} = \frac{1}{m}\sum_j \int_0^1 \big[(1-\epsilon)F^j_{B_1,A_1} + \epsilon F^j_{A_2,B_2}\big] dF^j_{A_1,B_1} + \lambda_{A_1,B_1}\left(X_{A,B} - \sum_j \int_0^1 x\, dF^j_{A_1,B_1}\right) \tag{9}$$

and

$$\mathcal{L}_{A_2,B_2} = \frac{1}{m}\sum_j \int_0^1 \big[(1-\epsilon)F^j_{B_2,A_2} + \epsilon F^j_{A_1,B_1}\big] dF^j_{A_2,B_2} + \lambda_{A_2,B_2}\left(X_{A,B} - \sum_j \int_0^1 x\, dF^j_{A_2,B_2}\right). \tag{10}$$

The resulting Euler–Lagrange equations must be solved taking into account the symmetry between players of the same type. They yield three different solutions, which in the following we label I, II, and III. Like in the two-player game, the index j can be dropped because of the mutual equivalence of the m items.

Using again the rescaled variables $u_{A,B} = (2X_B/m)^{-1} x_{A,B}$, the total resource ratio $\alpha = X_A/X_B$ (≤ 1), and the function $U(u; u_1, u_2)$ defined just below Eq. (8), the marginal probability distributions in solution I read

$$P^{\mathrm{I}}_A(u_A) = \frac{U[u_A; 0, \epsilon(1-\alpha)]}{\epsilon[1-\epsilon(1-\alpha)]} + \frac{[\alpha-\epsilon(1-\alpha)]U[u_A; \epsilon(1-\alpha), 1-\epsilon(1-\alpha)]}{[1-2\epsilon(1-\alpha)][1-\epsilon(1-\alpha)]} \tag{11}$$

and

$$P^{\mathrm{I}}_B(u_B) = \frac{U[u_B; \epsilon(1-\alpha), 1-\epsilon(1-\alpha)]}{1-2\epsilon(1-\alpha)}. \tag{12}$$

This solution is non-negative for $\alpha-\epsilon(1-\alpha)\ge 0$. The average payoffs are $\langle\pi_A\rangle = \alpha/2[1-\epsilon(1-\alpha)]$ and $\langle\pi_B\rangle = 1-\alpha/2[1-\epsilon(1-\alpha)]$.

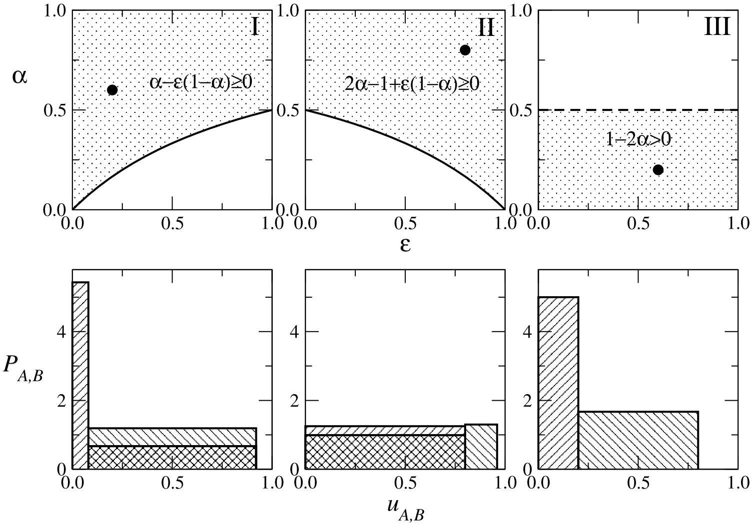

Fig. 2 *Upper panels*: *Shaded zones* show the domains in the parameter space (ϵ, α) where the solutions I, II, and III, given by Eqs. (11) to (15), are respectively well defined. *Lower panels*: Examples of the corresponding marginal probability distributions $P_A(u_A)$ (/-*dashed*) and $P_B(u_B)$ (\\-*dashed*) for each solution, at the points of parameter space marked with a *dot* in the *upper panels*

Solution II is given by

$$P_A^{\mathrm{II}}(u_A) = \frac{U(u_A; 0, \alpha)}{\alpha} \tag{13}$$

and

$$P_B^{\mathrm{II}}(u_B) = \frac{[2\alpha - 1 + \epsilon(1-\alpha)]U(u_B; 0, \alpha)}{\alpha[\alpha + \epsilon(1-\alpha)]} + \frac{U[u_B; \alpha, \alpha + \epsilon(1-\alpha)]}{\epsilon[\alpha + \epsilon(1-\alpha)]}. \tag{14}$$

It is non-negative for $2\alpha - 1 + \epsilon(1-\alpha) \geq 0$, and the average payoffs are $\langle \pi_A \rangle = 1 - 1/2[\alpha + \epsilon(1-\alpha)]$ and $\langle \pi_B \rangle = 1/2[\alpha + \epsilon(1-\alpha)]$.

Finally, solution III is

$$P_A^{\mathrm{III}}(u_A) = \frac{U(u_A; 0, \alpha)}{\alpha}, \qquad P_B^{\mathrm{III}}(u_B) = \frac{U(u_B; \alpha, 1-\alpha)}{1 - 2\alpha}. \tag{15}$$

It is positive and well-defined for $\alpha < 1/2$, and $\langle \pi_A \rangle = \epsilon/2$, $\langle \pi_B \rangle = 1 - \epsilon/2$.

The distributions in the three solutions are piecewise constant. The upper panels in Fig. 2 show the domains in the parameter space (ϵ, α) where each solution is well defined. Note that these domains overlap with each other over respectably large zones. Even more, just below $\alpha = 1/2$ and for moderate values of ϵ there is a zone where the three solutions coexist. The lower panels of the same figure show examples of the three solutions for selected parameter sets.

In the three solutions, the renormalized resource variables $u_{A,B}$ are divided into two ranges—the same for the two variables—within which the two marginal probability distributions $P_{A,B}$ are uniform. Solution I is characterized by the fact that A-players allocate both low and high resources, while B-players allocate high resources only. For a given A-player, a low resource allocation is convenient to maximize the payoff when competing against a player of the same type, while higher resources are necessary against B-players. It is interesting to note that solution I reproduces the distributions obtained for two players in Sect. 2.1

in the limits $\epsilon \to 0$ (when the opponents are players of different types only) and for $\alpha \to 1$ (when the two types of player become indistinguishable from each other), but not for $\epsilon \to 1$ (when the opponents are always of the same type).

Solution II is, in a sense, symmetric to solution I. In this case, A-players always choose to allocate low resources. B-players, on the other hand, allocate both low and high resources. The former optimize their payoff when playing against A-players, while the latter are necessary when playing against each other. This solution approaches the distributions obtained in Sect. 2.1 for $\epsilon \to 1$, and for $\alpha \to 1$, but not for $\epsilon \to 0$.

Finally, in solution III A- and B-players respectively choose to allocate low and high resources only. A-players thus resign winning any item in the competition with B-players but, on the other hand, maximize their payoffs when playing against each other. The only situation in which solution III reduces to the distributions obtained in Sect. 2.1 is for the trivial limit $\alpha \to 0$, when A-players have no resources at all.

Hence, the solutions stand for three different kinds of global strategies, representing the convenience of allocating either low or high resources in the competition against players with the same or with different (higher or lower) total resources, in the aim of maximizing the expected individual payoff. Whether one given strategy is more efficient than another depends on both the probability of confronting an opponent of either type—which is determined by the parameter ϵ—and the relation between the respective total resources—given by α. The question thus arises, for a given parameter set, as to which strategy prevails in the regions of the parameter space where two or three strategies coexist.

We choose to deal with such question by introducing a dynamical model for the Colonel Blotto game, where individual strategies evolve by imitation of the most successful players and progressively approach a stage where payoffs are maximized. This model, which is defined and studied numerically in Sect. 3, has the conceptual advantage of proposing an evolutionary picture that goes beyond the static representation of optimal strategies as Nash equilibria. At the same time, since numerical simulations make it possible to deal with large populations, we are able to validate the four-player representation used in this section.

3 Evolution of Strategies: Dynamical Rules and Numerical Simulations

In our evolutionary model of a population playing the Colonel Blotto game, we consider a set of N players, N_A of type A and N_B of type B. As before, the respective total resources are X_A and X_B for each player. The strategy of player n is characterized by the probability distribution $P_n(x,t)$ of allocating x resources at any given item at time t. The game is played successively by randomly chosen pairs of opponents and, as explained below, the players' strategies evolve in response to the resulting payoffs. As in Sect. 2.2, ϵ and $1-\epsilon$ are the respective probabilities that the opponent of any given player is of the same type and of different type.

In principle, the game dynamics occurs over two well-differentiated time scales. In the short term, many matches are played and information about the relative success of players with different strategies is stored. In a larger scale, this information is used to implement changes in the individual strategies, in an attempt to improve the own performance and maximize payoff. As is customary in evolutionary game theory [14], we bypass the shorter time scale by making the evolution of strategies to depend on the comparison between the expected payoffs of different players. This simplification amounts to assuming that many matches are played before strategies have the chance to change, so that any player is exposed to a representative sample of all the strategies over the population.

The dynamical rules for strategy evolution are defined at discrete time steps. At each step t, the average payoff expected for player n is calculated as

$$\langle \pi_n(t) \rangle = \frac{1-\epsilon}{N_{B,A}} \sum_{n_{B,A}} \int_0^{X_{A,B}} dx\, P_n(x,t) \int_0^x dx' P_{n_{B,A}}(x',t) + \frac{\epsilon}{N_{A,B}} \sum_{n_{A,B}} \int_0^{X_{B,A}} dx\, P_n(x,t) \int_0^x dx' P_{n_{A,B}}(x',t), \qquad (16)$$

where the first and the second subindex hold when n is of type A and B, respectively. The summation index $n_{A,B}$ runs over the players of either type, but skips player n.

Then, for each player n, a second player n' of the same type as n is chosen at random from the whole population and their expected payoffs are compared. If $\langle \pi_n(t) \rangle < \langle \pi_{n'}(t) \rangle$, player n adopts the strategy of n', i.e. $P_n(x,t)$ is replaced by $P_{n'}(x,t)$. If, on the other hand, $\langle \pi_n(t) \rangle \geq \langle \pi_{n'}(t) \rangle$, n does not change the strategy. The comparison is done, on the average, once per player per time step. This is an event of errorless imitation or, in biological terms, of inheritance without mutations.

Irrespectively of whether imitation has occurred or not—in order to have genuine evolution, by allowing the population to explore the space of all possible strategies—the strategy of player n is additionally submitted to a process of mutation. This consists in summing to the probability distribution $P_n(x,t)$ a suitably chosen perturbation $\delta P_n(x,t)$. In order to maintain probability normalization and not to modify the amount of total resources of player n, the perturbation must satisfy

$$\int_0^{X_{A,B}} dx \delta P_n(x,t) = 0, \qquad \int_0^{X_{A,B}} dx x \delta P_n(x,t) = 0. \qquad (17)$$

Moreover, it must be insured that the perturbed distribution is non-negative, $P_n(x,t) + \delta P_n(x,t) \geq 0$. How this is carried out in numerical simulations is explained below.

Once an initial condition—given by the set of probability distributions $P_n(x,0)$ for $n = 1, \dots, N$—has been specified, imitation and mutation are successively applied at each time step over the whole population. In order to implement these rules numerically, it is necessary to reduce the allocated resource variable x to a set of discrete values. To this end, we divide the interval $(0, 2X_B/m)$ into K equal parts, and take the points $x_k = 2X_B(k-1/2)/mK$ for $k = 1, 2, \dots$ as the possible values of the variable. The integrals in Eq. (16) are calculated using this discretization. Working with discretized variables, however, there is now a finite probability that the resources allocated by two opponents to a given item are exactly the same. Average payoffs are therefore computed assuming that, in such tied matches, the item is won by either player with probability $1/2$.

Over the discretized variable, we have applied the perturbation δP_n using two algorithms. In the first one, an amount $2q$ is subtracted from the probability $P_n(x_k,t)$ at a randomly chosen value x_k, and amounts q are added to the probabilities at the two nearest values, $P_n(x_{k+1},t)$ and $P_n(x_{k-1},t)$. In the second, the inverse process is performed: amounts q are subtracted from $P_n(x_{k+1},t)$ and $P_n(x_{k-1},t)$, and an amount $2q$ is added to $P_n(x_k,t)$. The two algorithms, which act only with the proviso that the non-negativity of probabilities is preserved, can straightforwardly be shown to satisfy the conditions of Eq. (17). Any of them is chosen with equal probability at every evolution step for each player, thus enhancing randomness in the process of strategy mutation.

We have performed extensive numerical simulations of the model, for various parameter sets (ϵ, α), with $N_A = N_B = 100$ ($N = 200$), $K = 40$, and $q = 1.25 \times 10^{-3}$. Several kinds of initial conditions for the probability distributions describing the strategy of each player

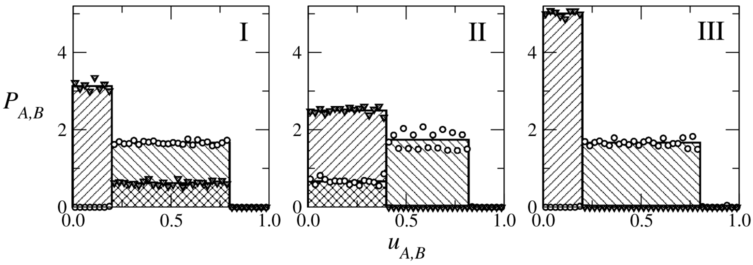

Fig. 3 Numerical results (*dots*) for the asymptotic probability distributions at three points of the parameter space (ϵ, α), from left to right, $(0.4, 0.5)$, $(0.7, 0.4)$, and $(0.6, 0.2)$. One of the three analytical distributions has been plotted for each parameter set using the same shading code as in Fig. 2, from left to right, solutions I, II, and III

were considered, including (i) multinomial distributions over $x_1, \ldots, x_K$; (ii) $P_n(x_k, 0)$ chosen from a uniform distribution for each n, and for each $k = 1, \ldots, K$; and (iii) distributions concentrated in a few intermediate values of x_k. In all cases, the distributions were properly normalized and satisfied condition (5), which insures the correct statistical representation of the total resources of each player. We also tried with a few initial conditions that included non-zero probabilities for values of x_k with $k > K$, finding however no qualitative differences with the other choices.

Since we were interested in a comparison between numerical simulations and the analytical results obtained in Sect. 2 for the optimal strategies we focused on the long-time behavior of our evolutionary model. In order to assess to which extent an asymptotic stage had been reached, we have evaluated the time dependence of a series of global quantities defined over the whole population, such as the average of the individual expected payoffs and the average mean square dispersion of the individual probability distributions. In all cases, and for all the initial conditions described above, we found that the system had attained a satisfactorily well-defined asymptotic stage after a time of about 10^5 steps.

Dots in Fig. 3 represent numerical results for the probability distributions of each kind of player, and three parameter sets (ϵ, α). Each numerical distribution was obtained as an average over the individual distributions of the players of each type, an over time. The time average was performed over 400 equally spaced time steps along a period of 1.6×10^5 steps, after discarding a transient of 2.4×10^5 steps from the initial condition. We verified that, for each pair (ϵ, α), the asymptotic distributions are independent of the initial condition.

Continuous lines in each panel of Fig. 3 represent one of the three solutions obtained in Sect. 2.2 for the corresponding parameter pair (ϵ, α)—from left to right, solutions I, II, and III. We find good agreement between the numerical results and the analytical solutions, indicating that the asymptotic distributions of our evolutionary model coincide with the optimal strategies obtained from joint maximization of the individual payoffs. At the same time, this coincidence with numerical results for a total population of $N = 200$ players validates the analytical four-player approximation of Sect. 2.2. We have verified that the agreement between analytical and numerical solutions occurs for a variety of parameter choices, covering the different behaviors of the system.

Figure 4 shows, in parameter space, various pairs (ϵ, α) for which one of the three optimal strategies found in Sect. 2.2 is obtained as the asymptotic solution of our evolutionary model. Diamonds represent parameter sets where the asymptotic solution coincides with solution III. This happens where the only optimal strategy is, precisely, solution III, i.e. below the curve $\alpha - \epsilon(1 - \alpha) = 0$ for $\epsilon < 1/2$ and below the curve $2\alpha - 1 + \epsilon(1 - \alpha) = 0$

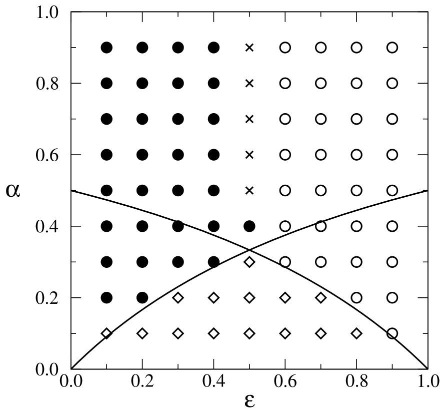

Fig. 4 Asymptotic solutions obtained numerically at different points of parameter space. *Full dots*: solution I; *empty dots*: solution II; *diamonds*: solution III. *Crosses* indicate parameter sets where none of the three solutions was approached within the computation time. The *curves* are the same as in the first two upper panels of Fig. 2, namely, $\alpha - \epsilon(1-\alpha) = 0$ and $2\alpha - 1 + \epsilon(1-\alpha) = 0$

for $\epsilon > 1/2$. In the zones where solution III coexists with the other solutions, on the other hand, it is never found to be asymptotically approached by the evolving strategies. Above the same curves, solutions I and II represent the asymptotic strategies for $\epsilon < 1/2$ (full dots) and $\epsilon > 1/2$ (empty dots), respectively. Note that each solution is obtained as the asymptotic strategy of the evolutionary model in the zone of parameter space where it correctly describes the limit cases $\epsilon \to 0, 1$ and $\alpha \to 0, 1$, as explained in Sect. 2.2.

For $\epsilon = 1/2$ and $\alpha \gtrsim 0.4$, we have not got well-defined asymptotic distributions within the above computation times. At the end of the simulations, probability distributions were still highly dependent on the initial condition and resembled neither solution I nor solution II. Since this zone is the boundary between the stability domains of those two solutions, we interpret this lack of convergence to a kind of critical slowing down associated with the instability transition of either solution.

It is interesting to realize that the asymptotic strategies obtained for any given pair (ϵ, α) correspond to the solution which minimizes the difference between the average expected payoff of the two types of players, $\langle \pi_B - \pi_A \rangle$. For each one of the three solutions, this quantity equals

$$\begin{aligned} \langle \pi_B - \pi_A \rangle_{\mathrm{I}} &= 1 - \alpha\big[1 - \epsilon(1-\alpha)\big]^{-1}, \\ \langle \pi_B - \pi_A \rangle_{\mathrm{II}} &= \big[\alpha + \epsilon(1-\alpha)\big]^{-1} - 1, \\ \langle \pi_B - \pi_A \rangle_{\mathrm{III}} &= 1 - \epsilon. \end{aligned} \tag{18}$$

Indeed, for fixed ϵ and α, the minimum among these three functions is obtained for the solution which represents the asymptotic probability distribution of our evolutionary model at that point of parameter space. Seen as a functional of the individual probability distributions, the average payoff difference $\langle \pi_B - \pi_A \rangle$ can be interpreted as a kind of non-equilibrium potential, whose minimization over the space of distributions yields the optimal strategies for a given parameter set. This non-equilibrium potential could be used to propose a thermodynamic-like picture for the evolution of strategies in this kind of game.

It might come as a surprise, however, that optimal strategies—which are obtained from a joint maximization of payoffs for all players—correspond to minimizing a function such as $\langle \pi_B - \pi_A \rangle$ which decreases for increasing π_A but *grows* with π_B. The solution to this seemingly paradoxical fact comes from the condition that, Colonel Blotto's being a constant-sum game, we always have $\langle \pi_B + \pi_A \rangle = 1$. Taking into account this condition, and the fact that $\langle \pi_B - \pi_A \rangle \geq 0$—because the total resources of B-players are larger or equal that those of A-players—it is straightforwardly verified that the minimization of the average payoff difference is equivalent to that of

$$\langle \pi_B - \pi_A \rangle^2 = 1 - 4\langle \pi_A \rangle \langle \pi_B \rangle. \tag{19}$$

This is in turn equivalent to *maximizing* the product $\langle \pi_A \rangle \langle \pi_B \rangle$, a symmetric function which grows with both payoffs.

4 Generalizations

Our evolutionary model for resource allocation strategies in the Colonel Blotto game admits to be extended in several directions. We have implemented some generalizations—briefly described in the following—which include (i) assigning different interaction probabilities to the two types of player; (ii) attributing an underlying interaction pattern, i.e. a "social structure," to the population; and (iii) relaxing the determinist nature of the imitation of strategies, thus adding a new stochastic ingredient to the evolution. In all cases, analytical solutions and numerical simulations disclose no qualitative differences with our previous findings, pointing to the robustness of the results.

4.1 Player-Dependent Interaction Probabilities

It may have come to the reader's mind at the very beginning of Sect. 2.2 that the probability of interaction of a given player with a player of the same type—the parameter ϵ of our model—needs not to be the same for both types of player. An evident generalization of the model is thus to introduce two parameters, ϵ_A and ϵ_B, respectively giving such probabilities for each type. The corresponding probabilities of confronting an opponent of the other type are $1 - \epsilon_A$ and $1 - \epsilon_B$. These two parameters arise naturally, for instance, if the probability of different matches are defined in terms of opponent pairs. In fact, if γ_A and γ_B are the probabilities that the game is played by two A-players and two B-players, respectively—while $1 - \gamma_A - \gamma_B$ is the probability for the players being of different types—we find $\epsilon_A = \gamma_A/(1 - \gamma_B)$ and $\epsilon_B = \gamma_B/(1 - \gamma_A)$, which are generally different. This case includes also the relevant situation where the two opponents are independently chosen from the whole population, but the number of players of each type is not the same. If the respective fractions of players are ρ_A and $\rho_B = 1 - \rho_A$, we find $\gamma_A = \rho_A^2$ and $\gamma_B = \rho_B^2$, and the parameters turn out to be $\epsilon_A = \rho_A/(2 - \rho_A)$ and $\epsilon_B = \rho_B/(2 - \rho_B)$.

The analytical problem arising from this generalization is essentially the same as in Sect. 2.2 and, correspondingly, its solution does not show major differences with respect to the case of $\epsilon_A = \epsilon_B$. As before, three piecewise constant solutions exist for each parameter set. Both the domains within which the three distributions are constant, and the regions of parameter space where the solutions are well defined, depend on ϵ_A, ϵ_B, and α. The asymptotic behavior of the model should now be numerically studied on a three-dimensional space. Otherwise, the solutions are qualitatively the same as in the simpler case of two parameters. Apart from the extra parameter, therefore, this extension does not contribute any substantial novelty to the model.

4.2 Playing the Game on Small-World Networks

Numerical simulations make it possible to introduce a second, important generalization, regarding the social structure of the population. Specifically, we have considered the possibility that each player can interact with a prescribed subpopulation only, which defines the player's "neighborhood." Such structure is conveniently implemented by means of a network, whose nodes represent the players and whose links join players who can confront each other in the game. This generalization allows, in principle, for the consideration of complex social structures and for situations where each player can only confront a small part of the whole population, which are realistic scenarios beyond the reach of analytical treatment.

In our simulations, we have used small-world networks built from one-dimensional ordered arrays with $N = 200$ and 2000 nodes, and with $z = 10$ neighbors per node (5 to each side), following Watts and Strogatz's rewiring algorithm [15]. We have assumed, as explained above, that opponents are always chosen within individual neighborhoods—so that the calculation of the expected payoff of each player only involves the strategies of the corresponding neighbors—but that imitation can occurs over the whole population. This adds the interesting ingredient that a strategy which is successful in a given neighborhood, with its specific number of players of each type, may be not convenient in other environments. Numerical results show that, indeed, there is a moderately high correlation between the asymptotic strategy approached by each player and the composition of the local neighborhood. We have evaluated Pearson's correlation coefficient r between the average individual payoff and the local fraction of players of each type and obtained values around $r \approx 0.2$. However, the simulations suggest that the probability distributions observed for long times are of the same type as those obtained from the analytical calculations. Moreover, results are not sensible to the degree of disorder of the small-world network, as measured by the rewiring parameter. We conclude that, at least for this kind of network, our previous results are robust under changes in the social structure of the population.

4.3 Non-deterministic Imitation

Finally, in order to add randomness to the evolution, strategy imitation between players can be generalized by admitting that—instead of deterministically depending on the comparison of average payoffs—the adoption of another player's strategy is controlled by a probability. In our generalization of imitation, we have considered that player n adopts the strategy of player n' with a probability $p(\Delta)$, which depends on the average payoff difference $\Delta = \langle \pi_{n'} \rangle - \langle \pi_n \rangle$. Specifically, as above, if this difference is negative imitation is rejected: $p(\Delta < 0) = 0$. For positive Δ, on the other hand, we have taken $p(\Delta > 0) = 1 - \exp(-w\Delta)$, so that imitation is more likely as Δ grows. The positive parameter w defines how rapidly the probability varies with Δ. The deterministic algorithm introduced in Sect. 2.2 corresponds to the limit $w \to \infty$. We have performed simulations for $w = 0.1$ and 1, and found no essential differences with the deterministic case, except for the expected slowing down of the overall dynamics as w—and, consequently, the imitation probability for any fixed Δ—decreases.

A generalization along a similar line would have been to admit that there is a non-zero probability that player n adopts the strategy of player n' even when this imitation implies a decrease of the payoff ($\Delta < 0$). A standard implementation of this procedure is the Metropolis algorithm [16], taking $p(\Delta < 0) = \exp(\Delta/T)$ and $p(\Delta > 0) = 1$. The "temperature" T measures the degree of randomness in the process. The imitation algorithm introduced of Sect. 2.2 corresponds to zero temperature. As in other statistical problems, the introduction of temperature could play a role preventing the system to become stuck on metastable equilibria. In our case, however, there is no evidence of the existence of such states.

5 Conclusion

Strategies of resource allocation are a basic ingredient of economic and financial management, and become increasingly crucial as the complexity of the involved systems and the pressure of competition grow. They also plays a role in social organization and in biological evolution. As it directly underlies the maximization of benefits for a given investment, resource allocation is a problem naturally adapted for treatment within game theory. The Colonel Blotto game, which we have studied in this paper, is just one of a class of models exploiting game theory in that direction.

Here, our main goal has been to frame the Colonel Blotto game within evolutionary game theory. The Colonel Blotto game considers the competition of two players with different total resources to be distributed among a set of items. Each item is won by the player that allocated higher resources to it. We have begun our analysis by treating, within the traditional theory, a generalized game where a given player can confront an opponent with either the same or different resources. This situation was conveniently represented by a four-player population, two of them with low resources and the other two with high resources. For this simplified version of a larger population we have been able to obtain optimal strategies by means of a Lagrangian maximization of payoffs. We obtained three solutions, partially coexisting in parameter space, describing strategies which favor disparate forms of resource distribution in the competition with players with equal and different resources. Each type of player can adopt two kinds of behavior. In one of them, common to both types, strategies are distributed over the whole range of resources. In the other behavior, in contrast, players with low and high resources respectively concentrate their strategies at the lower and higher ends of the same range.

We then moved to the definition of evolutionary rules for strategies, based—much like biological evolution—on the combined action of strategy imitation and variation by random fluctuations. In this dynamical framework, optimal strategies are expected to be found at the asymptotic long-time stages of the evolving population. We have performed extensive numerical simulations in populations of a few hundred players and found that, effectively, the asymptotic strategies coincide with those obtained by Lagrangian maximization. These results validate the four-player representation used in our analytical calculations. More importantly, numerical simulations make it possible to identify, in the zones of parameter space where more than one analytical solution coexist, which of them is actually adopted at long times. Comparing with the expected average payoffs of each solution, we concluded that the preferred strategy is the one that maximizes the product of payoffs of the two types of player. This product, therefore, acts as a kind of non-equilibrium potential for the evolutionary process.

Finally, we have verified that our main conclusions are robust against a series of generalizations of the model, both in the game rules and in the evolutionary dynamics. They included relaxing some of the symmetries between players in the original version, considering a population with heterogeneous local structure—as represented by a small-world network—and adding extra stochastic ingredients to the imitation of strategies. In all cases we observed no significant changes in the overall behavior of our system.

Previous work formulated and analyzed the Colonel Blotto game within traditional game theory, and thus was limited to providing a static scenario of optimal resource allocation. Evolutionary game theory, on the other hand, sets up a dynamical framework with the participation of a whole population of agents, where optimal strategies are the outcome of evolution itself. As such, it provides a more realistic picture of most socio-economical (or biological) systems, whose nature is inherently dynamical.

Acknowledgements We acknowledge financial support from CONICET (PIP 112-200801-76) and ANPCyT (PICT 2011-545), Argentina.

References

1. Borel, É.: La théorie du jeu et les équations intégrales à noyau symétrique gauche. C. R. Hebd. Séances Acad. Sci. **173**, 1304–1308 (1921)
2. Borel, É., Ville, J.: Application de la théorie des probabilitiés aux jeux de hasard. Gauthier-Villars, Paris (1938)
3. Gross, O., Wagner, R.: A continuous Colonel Blotto game, RM-408. RAND Corporation, Santa Monica (1950).
4. Baye, M.R., Kovenock, D., de Vries, C.G.: The all-pay auction with complete information. Econ. Theory **8**, 291–305 (1996)
5. Myerson, R.B.: Incentives to cultivate minorities under alternative electoral systems. Am. Polit. Sci. Rev. **87**, 856–869 (1993)
6. Laslier, J.F., Picard, N.: Distributive politics and electoral competition. J. Econ. Theory **103**, 106–130 (2002)
7. Sahuguet, N., Persico, N.: Campaign spending regulation in a model of redistributive politics. Econ. Theory **28**, 95–124 (2006)
8. Roberson, B.: The Colonel Blotto game. Econ. Theory **29**, 1–24 (2006)
9. Golman, R., Page, S.E.: General Blotto: games of allocative strategic mismatch. Public Choice **138**, 279–299 (2009)
10. Kvasov, D.: Contests with limited resources. J. Econ. Theory **127**, 738–748 (2007)
11. Hart, S.: Discrete Colonel Blotto and general Lotto games. Int. J. Game Theory **36**, 441–460 (2008)
12. Kovenock, D., Roberson, B.: Conflicts with multiple battlefields. Working Paper Series No. 3165. CESifo, Munich (2010)
13. Maynard Smith, J.: Evolution and the Theory of Games. Cambridge University Press, Cambridge (1982)
14. Hofbauer, J., Sigmund, K.: Evolutionary Games and Population Dynamics. Cambridge University Press, Cambridge (1998)
15. Watts, D.J., Strogatz, S.H.: Collective dynamics of ‘small-world’ networks. Nature **393**, 440–442 (1998)
16. Newman, M.E.J., Barkema, G.T.: Monte Carlo Methods in Statistical Physics. Oxford University Press, Oxford (1999)