



An Application of Evolutionary Game Theory to Social Dilemmas: The Traveler's Dilemma and the Minimum Effort Coordination Game

Citation

Iyer, Swami, Joshua Reyes, and Timothy Killingback. 2014. "An Application of Evolutionary Game Theory to Social Dilemmas: The Traveler's Dilemma and the Minimum Effort Coordination Game." PLoS ONE 9 (4): e93988. doi:10.1371/journal.pone.0093988. <http://dx.doi.org/10.1371/journal.pone.0093988>

Published version

<https://doi.org/10.1371/journal.pone.0093988>

Link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:12153046>

Terms of use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material (LAA), as set forth at

<https://harvardwiki.atlassian.net/wiki/external/NGY5NDE4ZjgzNTc5NDQzMGIzZWZhMGFIOWI2M2EwYTg>

Accessibility

<https://accessibility.huit.harvard.edu/digital-accessibility-policy>

Share Your Story

The Harvard community has made this article openly available. Please share how this access benefits you. [Submit a story](#)



An Application of Evolutionary Game Theory to Social Dilemmas: The Traveler's Dilemma and the Minimum Effort Coordination Game

Swami Iyer¹, Joshua Reyes², Timothy Killingback^{3*}

1 Computer Science Department, University of Massachusetts, Boston, Massachusetts, United States of America, **2** Department of Systems Biology, Harvard Medical School, Boston, Massachusetts, United States of America, **3** Mathematics Department, University of Massachusetts, Boston, Massachusetts, United States of America

Abstract

The Traveler's Dilemma game and the Minimum Effort Coordination game are two social dilemmas that have attracted considerable attention due to the fact that the predictions of classical game theory are at odds with the results found when the games are studied experimentally. Moreover, a direct application of deterministic evolutionary game theory, as embodied in the replicator dynamics, to these games does not explain the observed behavior. In this work, we formulate natural variants of these two games as smoothed continuous-strategy games. We study the evolutionary dynamics of these continuous-strategy games, both analytically and through agent-based simulations, and show that the behavior predicted theoretically is in accord with that observed experimentally. Thus, these variants of the Traveler's Dilemma and the Minimum Effort Coordination games provide a simple resolution of the paradoxical behavior associated with the original games.

Citation: Iyer S, Reyes J, Killingback T (2014) An Application of Evolutionary Game Theory to Social Dilemmas: The Traveler's Dilemma and the Minimum Effort Coordination Game. PLoS ONE 9(4): e93988. doi:10.1371/journal.pone.0093988

Editor: Matjaž Perc, University of Maribor, Slovenia

Received: January 13, 2014; **Accepted:** March 11, 2014; **Published:** April 7, 2014

Copyright: © 2014 Iyer et al. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Funding: The authors have no support or funding to report.

Competing Interests: The authors have declared that no competing interests exist.

* E-mail: timothy.killingback@umb.edu

Introduction

Social dilemmas embody the tension between individual self-interest and the common good that is inherent in many important situations in the real world. In a social dilemma, individually reasonable behavior results in a situation in which all individuals are less well off than they could otherwise have been [1]. Social dilemmas underlie many of the most fundamental and intractable problems in the biological and social sciences, such as the evolution of cooperation [2] and the efficient use of limited shared resources [3]. From a more formal point of view, a social dilemma can be modeled as a game in which there exists at least one Hicks inefficient Nash equilibrium. It is Hicks inefficient (i.e. socially inefficient) in that there is at least one other outcome in which all individuals would be better off, and since it is a Nash equilibrium there is no incentive for any individual to change their behavior [1]. Examples of 2-person games that are social dilemmas include: The Prisoner's Dilemma [2,4,5], the Snowdrift game (also known as the Chicken or Hawk-Dove game) [6], and the Stag-Hunt game (also known as the Assurance game) [7]. Multi-person social dilemmas include, the Public Goods game [8,9] and the Tragedy of the Commons [3]. While the original game theory formalizations of these social dilemmas typically involved discrete-strategy games, more recently continuous-strategy versions of the Prisoner's Dilemma game [10], the Snowdrift game [11], the Tragedy of the Commons [12] game, and the Public Goods game [13] have been formulated and studied.

One particularly fascinating class of social dilemmas are those for which the predictions of game theory appear to be inconsistent

with the behavior observed when the games are played experimentally [14]. The Traveler's Dilemma (TD) game [15–17] and the Minimum Effort Coordination (MEC) game [17,18] are two celebrated examples of such games.

These games challenge the notion that the rational solution to a game, as embodied in the concept of the Nash equilibrium, accurately describes the behavior of humans engaged in these social dilemmas. In the TD game there exists a unique, pure strategy, Nash equilibrium that is undesirable for all concerned. In contrast to the TD game, in the MEC game every pure strategy is a Nash equilibrium, and thus the rational solution concept lacks any prescriptive or predictive power.

The TD game is conventionally introduced through a story of the following form. Two travelers, on their return journey from an exotic country, find that their luggage containing identical souvenirs has been lost by the airline. The officer in the claims department puts them in separate rooms, hands each of them a claims form, and tells them that they can claim any integer amount between R and M (R and M are assumed to be positive integers with $M > R$). He also informs them that if they both ask for the same amount, they will be paid that amount, and if they ask for different amounts, each will be reimbursed at the lower value, but with a penalty R deducted from the higher claimant (who is assumed to have lied) and given to the lower claimant (as a reward for being honest). Thus, the TD game is a 2-person game with the discrete strategy set $\mathcal{S} = \{R, R+1, R+2, \dots, M\}$, and payoff to an i -claimant against a j -claimant is defined by

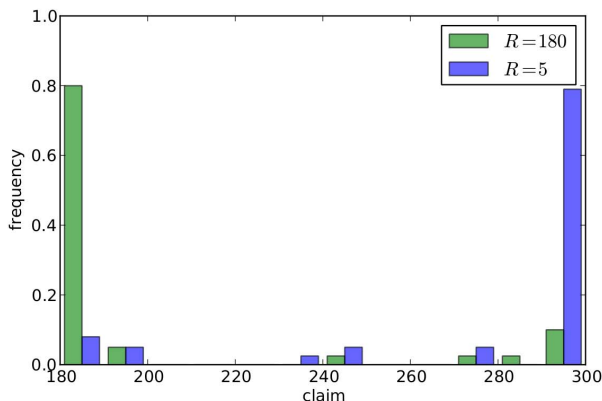
$$P(i,j) = \begin{cases} i+R & \text{if } i < j \\ i & \text{if } i = j \\ j-R & \text{if } i > j. \end{cases} \quad (1)$$

In this context a Nash equilibrium is a pair of claims, such that, if each claim is known to the other traveler then neither has reason to revise their claim. For $R > 1$, there is an incentive for each traveler to undercut any common claim. Using backward induction, it is not hard to see that the travelers should each claim the amount R , i.e., (R, R) is a unique Nash equilibrium for the TD game. Thus, the unique Nash equilibrium of the TD game is the paradoxical outcome in which both travelers claim the lowest possible amount. We note, in particular, that the Nash equilibrium is independent of the reward/punishment parameter R . However, as intuition would suggest, this is not how individuals actually play this game [17,19–22]. For instance, [22] found the following results when they played the game with 50 subjects (25 pairs). The subjects could make claims between 180 and 300, in two treatments, one with $R = 180$ and another with $R = 5$. The results are shown in Figure 1(a). In the high- R treatment, close to 80 percent of all the subjects chose the Nash equilibrium strategy, with an average claim of 201. However, in the low- R treatment, roughly the same fraction chose the highest possible claim, with an average value of 280. Since the unique Nash equilibrium prediction is independent of the parameter R , classical game theory is unable to explain the most salient feature of these experimental results, namely, the effect of the reward/punishment parameter R on average claim levels.

The MEC game has a somewhat similar flavor to the TD game, in that the payoffs are again determined by the minimum of two actions [17,18,20]. In this game the players choose integer effort levels between 1 and M (where M is assumed to be an integer greater than 1), and a player's payoff is given by the minimum of the two effort levels minus the cost of the player's own effort. Therefore, the MEC game is a 2-person game with the discrete strategy set $S = \{1, 2, \dots, M\}$, and payoff to an i -strategist against a j -strategist defined by

$$P(i,j) = \min(i,j) - c \cdot i, \quad (2)$$

(a)



(b)

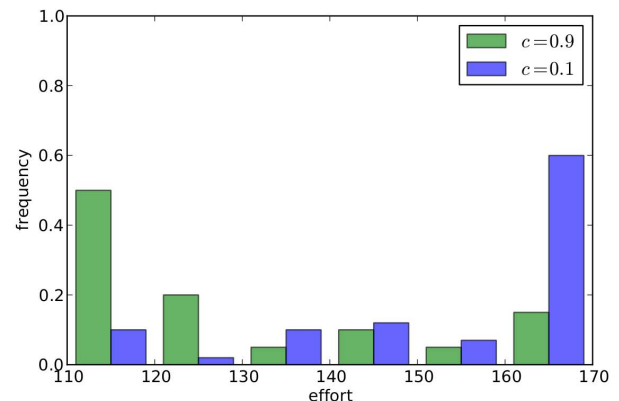


Figure 1. Results from playing the TD and MEC games with human subjects, adapted from [17]. (a) TD game: individuals make higher claims when the reward/punishment parameter R is low, and make lower claims when R is high. (b) MEC game: individuals expend more effort when the effort cost c is low, and less effort when c is high. doi:10.1371/journal.pone.0093988.g001

species of communally nesting birds, the Greater Ani [25,26]. In this species, if two females share a nest then each female chooses a time to change from ejecting eggs from the nest to laying eggs. If both select an early time (which corresponds to a large claim), then both obtain a large payoff since they can both successfully lay many eggs. If, however, one chooses to wait, then she can eject the other birds already-laid eggs and obtain an even greater payoff, while at the same time inflicting a loss on the other bird [24]. Thus, this situation has the structure of the TD game.

The MEC game also has both theoretical and practical importance. It is theoretically significant because it starkly illustrates the lack of prescriptive or predictive power inherent in the notion of Nash equilibrium. In addition, it has considerable practical import since many interesting and important real-world situations can be modeled by MEC games [27].

As an example of how MEC games naturally arise let us consider two companies, denoted by A and B, each of which manufacture a critical component of a jointly produced product. Let us suppose that company A makes widgets, while company B makes grommets. The final product, containing both a widget and a grommet, is sold jointly, with the revenues being equally split between the two companies. Each company can choose the amount of effort to expend in producing its component, with higher effort levels resulting in components of higher quality. We shall assume that the performance of the final product, and thus also the revenues obtained from sales of the product, is limited by whichever of the two components has the lower quality. Therefore, the profits obtained by a given company from sales of the product may be expressed as the minimum of the efforts expended by each company to produce widgets or grommets, respectively, minus the cost associated with the companies' own effort. Thus, such a situation can be modeled by a MEC game.

The importance of the TD and MEC games, both theoretically and in practice, is clear from the above comments. Therefore, obtaining a satisfactory understanding of the dynamics of these games is of considerable significance, and it is the purpose of this paper to contribute to such an understanding.

A number of different theoretical approaches have been investigated as possible explanations of the behavior found empirically in the TD and MEC games. In one approach, stochastic learning models [17,20,28] have been proposed to explain the anomalous behavior observed in the TD and MEC games. A quite different approach, using stochastic evolutionary dynamics in finite populations [24], has also been investigated as a means of resolving the paradoxical features of the TD game. Other theoretical approaches to explaining the behavior of the TD game have been studied in [29–32]. Other approaches which explore how errors in a game may lead to deviations from Nash equilibrium play include [33–37].

Here we propose an alternative, considerably simpler, theoretical framework to explain the evolutionary dynamics of both the TD and MEC games, which accounts for the empirically observed behavior. Our method applies to a wide class of games that includes both the TD and MEC games. We first observe that while the TD and MEC games were originally formulated as discrete strategy games it is natural to consider variants of these games in which the strategies are continuously variable. In this paper we define these continuous-strategy variants of the TD and MEC games, which suffer from the same paradoxical behavior as the original discrete-strategy games, and it is these continuous-strategy games that form the starting point for our approach to understanding the evolutionary dynamics of the TD and MEC games.

The continuous-strategy forms of the TD and MEC games are two examples of a large class of continuous-strategy games which have a discontinuous payoff function — other important examples of games from this class include the Bertrand Duopoly game [38] and the War of Attrition game [6]. For such games the role played by errors is potentially important. To be precise, let us consider the class of continuous-strategy games with payoff function given by

$$P(x,y) = \begin{cases} \phi(x,y) & \text{if } x < y \\ \frac{1}{2}[\phi(x,y) + \psi(x,y)] & \text{if } x = y \\ \psi(x,y) & \text{if } x > y, \end{cases} \quad (3)$$

where $\phi(x,y)$ and $\psi(x,y)$ are affine functions, and the strategies $x,y \in \mathbb{R}_+$. This class of games includes the continuous-strategy variants of both the TD and MEC games. Errors in the observation of an opponents strategy or in the implementation of ones own strategy will result in the expected payoff to an x -strategist against a y -strategist in such a game being given by a function of the form

$$P(x,y) = \phi(x,y)\theta(y-x) + \psi(x,y)\theta(x-y), \quad (4)$$

where $\theta(y-x)$ defines the probability that in an interaction between an x -strategist and a y -strategist errors lead to the x -strategist receiving the payoff $\phi(x,y)$, and $\theta(x-y)$ defines the probability that in such an interaction errors result in the x -strategist receiving the payoff $\psi(x,y)$. We observe that the probability functions $\theta(x-y)$ and $\theta(y-x)$ are necessarily complementary in the sense that $\theta(x-y) + \theta(y-x) = 1$, from

which it follows that $\theta(0) = \frac{1}{2}$. The probability functions $\theta(x-y)$ and $\theta(y-x)$ are determined by the statistical distribution of errors in the game. Here, for simplicity, we shall assume that the function θ is *smooth*. Thus, the expected payoff (4) defines a *smoothing* of the original discontinuous payoff function (3). We shall often refer to the function θ as the smoothing function. In the limit in which the smoothing function tends to the Heaviside step function the expected payoff (4) approaches the payoff (3).

Since the purpose of this paper is to study the evolutionary behavior of the TD and MEC games it is necessary to define a suitable evolutionary dynamics for the class of games that we are considering. Since the effect of errors in the game is to give an expected payoff which is a smooth function, the simplest choice of dynamics is the standard deterministic adaptive dynamics [39–41] of the smoothed payoff function. In our approach the effects of errors in the game is encoded in the smoothing function θ . Here we study the evolutionary dynamics of the smoothed versions of the TD and MEC game, for an arbitrary smoothing function, and show that the results obtained are consistent with those found empirically.

The approach to studying the evolutionary dynamics of the TD and MEC games that we follow in this paper has some advantages compared to stochastic learning models [17,20,28] or stochastic evolutionary dynamics [24]. In the case of stochastic learning models, the evolutionary dynamics is governed by the Fokker-Planck equation, which is a nonlinear partial differential equation that cannot be solved analytically [17, 20, 28]. The equilibrium solutions of the evolutionary dynamics are given by the solutions to a suitable differential equation. However, it is difficult to determine analytically whether or not the equilibrium solutions are stable. In fact, the equilibrium solutions to the stochastic learning models [17,20,28] for both the TD and MEC games have

not been shown to be stable, and thus, it is unclear whether or not they are attractors of the evolutionary dynamics.

Stochastic evolutionary dynamics [24] represents an interesting alternative approach to understanding the dynamics of many evolutionary processes. In this case, the evolutionary dynamics is governed by a stochastic process, and the theory is mathematically well-developed. However, there are certain restrictions that apply to the theory. Perhaps the most important is that analytic results can only be obtained in the limit that the selection strength tends to zero. This limit corresponds to the assumption that the contribution to the total payoff that comes from the game interactions is very small. It is not clear that this assumption is realistic when it comes to understanding the behavior of the TD and MEC games. If the assumption of weak selection is not made then it is impossible to obtain analytic results, although numerical simulations can yield results for stronger selection strengths. It is also worth noting that while the stochastic process underlying stochastic evolutionary dynamics is mathematically well-developed, it is rather subtle, and this can present a challenge when applying the method to new problems. In fact, stochastic evolutionary dynamics has not yet been applied to study the MEC game, although we conjecture that such an application will yield results consistent with the behavior observed in the game.

A key advantage of the method that we propose here is that the evolutionary dynamics is much easier to study than for either stochastic learning models [17,20,28] or stochastic evolutionary dynamics [24]. In particular, it is straightforward to completely determine the evolutionary attractors for both the TD and MEC games. A second advantage of our method is that it applies directly to a wide variety of continuous-strategy games with discontinuous payoff functions, including the Bertrand Duopoly model [38] and the War of Attrition game [6]. In certain cases, which are considered at greater length in the Discussion, it can be shown using our methods that complex evolutionary dynamics, such as evolutionary branching, can occur in such games.

The rest of the paper is organized as follows. In the Models section, we define continuous-strategy versions of the TD and MEC games, and also introduce the key notion of smoothed versions of these games. In the Analysis section, we analyze the evolution of strategies in the smoothed games in randomly-interacting populations using adaptive dynamics, and in addition formulate an agent-based model of the evolutionary dynamics of these games in populations with structures described by an arbitrary graph (i.e., network). In the Results section, we present the results of simulations using this agent-based model for the evolutionary dynamics of the smoothed TD and MEC games, both in well-mixed populations and in populations described by complex networks. Finally, in Discussion section, we provide a brief discussion of our work and draw some conclusions.

Models

The strategies in the TD and the MEC games are the claim levels and the effort levels, respectively. Typically these games are taken to have a discrete set of strategies. However, it is in many ways more natural to view the claim levels and effort levels in the two games as being continuously variable, and thus to consider variants of these games defined for continuous strategies. The continuous forms of the TD and MEC games are examples of a broad class of continuous-strategy 2-person games with payoff functions given by

$$P(x,y) = \begin{cases} \phi(x,y) & \text{if } x < y \\ \frac{1}{2}[\phi(x,y) + \psi(x,y)] & \text{if } x = y \\ \psi(x,y) & \text{if } x > y, \end{cases} \tag{5}$$

for affine functions $\phi(x,y)$ and $\psi(x,y)$, and strategies $x,y \in \mathbb{R}_+$. We may write $P(x,y)$ more succinctly with the aid of the Heaviside step function $\Theta(u)$

$$\Theta(u) = \begin{cases} 0 & \text{if } u < 0 \\ \frac{1}{2} & \text{if } u = 0 \\ 1 & \text{if } u > 0, \end{cases} \tag{6}$$

as

$$P(x,y) = \phi(x,y)\Theta(y-x) + \psi(x,y)\Theta(x-y). \tag{7}$$

Games of this form have discontinuous payoff functions. Such a discontinuous payoff function is only possible in an idealized world free from all errors. In reality, errors in the perception and implementation of actions in the game will have the effect of replacing the discontinuous payoff function with a smoothed approximation, representing the expected payoff. We now define such a variant of the game in which the discontinuity in the payoff function is removed by a smoothing procedure. To accomplish this we introduce a 1-parameter family of smoothing functions $\theta_k(u)$. The functions $\theta_k(u)$ are assumed to be smooth, non-decreasing functions of u , with $\lim_{u \rightarrow +\infty} \theta_k(u) = 1$, $\lim_{u \rightarrow -\infty} \theta_k(u) = 0$, and $\theta_k(0) = \frac{1}{2}$. Furthermore, we assume that $\theta_k(u) \rightarrow \Theta(u)$ as $k \rightarrow \infty$. We will refer to the parameter k as the smoothing parameter.

To obtain the smoothed version of the game defined by (7) we simply replace $\Theta(u)$ in the payoff function with its smooth approximation $\theta_k(u)$. Thus, the payoff function of the smoothed game is given by

$$P_k(x,y) = \phi(x,y)\theta_k(y-x) + \psi(x,y)\theta_k(x-y). \tag{8}$$

We note that for sufficiently large values of k the smoothed game approximates the original game arbitrarily well.

A convenient 1-parameter family of smoothing functions is given by

$$\theta_k(x) = \frac{1}{1 + e^{-kx}}, \tag{9}$$

and we shall use this family when explicit smoothing functions are required.

Traveler's Dilemma Game

The claims made by individuals in the TD game represent their strategies. If x and y denote the strategies used by two individuals playing the continuous version of the game, and if R denotes the reward/punishment parameter (where $x,y,R \in \mathbb{R}_+$), then the payoff to the x -strategist is given by

$$P(x,y) = (x + R)\Theta(y-x) + (y - R)\Theta(x-y). \tag{10}$$

We now define a variant of the TD game defined by (10), in which the discontinuity in the payoff function is removed by the

smoothing procedure described above. To obtain the smoothed TD game we simply replace $\Theta(x)$ in the payoff function with its smooth approximation $\theta_k(x)$. We therefore have that the payoff function for the smoothed TD game is given by

$$P_k(x,y) = (x + R)\theta_k(y - x) + (y - R)\theta_k(x - y). \quad (11)$$

We shall assume, without loss of generality, that the strategy space in the smoothed TD game is the interval $[0,1]$, and also that the reward/punishment parameter $R \in (0,1)$. With the payoff function defined by (11), the smoothed TD game represents a natural variant of the original TD game.

Minimum Effort Coordination Game

In the MEC game, the effort levels of the individuals represent their strategies. In the continuous version of the MEC game, if the strategies of two individuals playing the game are x and y (where $x, y \in \mathbb{R}_+$) and $c \in (0,1)$ is the effort cost, then the payoff $P(x,y)$ to the x -strategist is given by

$$P(x,y) = \min(x,y) - c \cdot x. \quad (12)$$

Using (12) allows us to write the payoff function for the MEC game as

$$P(x,y) = (x - cx)\Theta(y - x) + (y - cx)\Theta(x - y). \quad (13)$$

Without loss of generality we can take the strategy space to be the unit interval (i.e. $x, y \in [0,1]$). Every strategy pair (x,x) is a Nash equilibrium in this game. The social dilemma embodied in this continuous-strategy game is clearly the same as for the original discrete MEC game: at any equilibrium both players obtain a payoff of $(1 - c)x$, thus all equilibria with the sole exception of the strategy pair $(1,1)$ are Hicks inefficient.

To obtain the smoothed MEC game we again replace $\Theta(x)$ in the payoff function (13) with its smooth approximation $\theta_k(x)$. The payoff function of the smoothed MEC game is therefore given by

$$P_k(x,y) = (x - cx)\theta_k(y - x) + (y - cx)\theta_k(x - y). \quad (14)$$

With the payoff function defined by (14), the smoothed MEC game represents a natural variant of the original MEC game.

Analysis

The dynamics of the smoothed TD and MEC games as formulated in the previous section can be analyzed in a well-mixed population using the deterministic framework of adaptive dynamics [11,39–41]. Consider a monomorphic population in which every individual adopts the same strategy, x . It follows from replicator dynamics that the growth rate of a rare mutant strategy, y , in the resident x population is $f_x(y) = P_k(y,x) - P_k(x,x)$, where $P_k(x,y)$ is the payoff to an x -strategist interacting with a y -strategist. The quantity $f_x(y)$ is referred to as the invasion fitness. The evolution of the strategy x is then governed by the selection gradient $D(x) = \frac{\partial f_x}{\partial y} \Big|_{y=x}$, and the adaptive dynamics of x is determined by the differential equation $\dot{x} = D(x)$.

Equilibrium points of the adaptive dynamics are called singular strategies and are solutions of $D(x^*) = 0$. If no such solutions exist, then the strategy x monotonically increases or decreases under evolution, depending on the sign of $D(x)$. If x^* exists, it is convergent stable and, hence an attractor for the adaptive

dynamics, if $\frac{dD}{dx} \Big|_{x=x^*} < 0$. If this equality is reversed, x^* is a repeller.

Initially, the population will converge to a convergent stable singular point x^* , but its subsequent evolutionary fate depends on whether x^* is a maximum or minimum of the invasion fitness $f_x(y)$. If x^* is a maximum, i.e., if $\frac{\partial^2 f_x}{\partial y^2} \Big|_{y=x^*} < 0$, then x^* is an evolutionarily stable strategy (ESS), representing an evolutionary end state in which all individuals adopt strategy x^* . If, however, $\frac{\partial^2 f_x}{\partial y^2} \Big|_{y=x^*} > 0$, then a population of x^* -strategists can be invaded by mutant strategies on either side of x^* . In this case the population undergoes evolutionary branching and splits into two distinct and diverging clusters of strategies.

The adaptive dynamics of smoothed games with payoff function defined by (8) may be analyzed as follows. The invasion fitness is given by

$$\begin{aligned} f_x(y) &= P_k(y,x) - P_k(x,x) \\ &= \phi(y,x)\theta_k(x - y) + \psi(y,x)\theta_k(y - x) - [\phi(x,x) + \psi(x,x)]\theta_k(0). \end{aligned} \quad (15)$$

Thus, the selection gradient $D(x)$ is given by

$$\begin{aligned} D(x) &= \frac{\partial f_x(y)}{\partial y} \Big|_{y=x} \\ &= \frac{\partial}{\partial y} [\phi(y,x)\theta_k(x - y) + \psi(y,x)\theta_k(y - x) - \phi(x,x)\theta_k(0) \\ &\quad - \psi(x,x)\theta_k(0)] \Big|_{y=x} \\ &= [\psi(x,x) - \phi(x,x)]\theta'_k(0) + \theta_k(0) \frac{\partial}{\partial y} [\phi(y,x) + \psi(y,x)] \Big|_{y=x} \\ &= [\psi(x,x) - \phi(x,x)]\theta'_k(0) + \frac{1}{2} \frac{\partial}{\partial y} [\phi(y,x) + \psi(y,x)] \Big|_{y=x}. \end{aligned} \quad (16)$$

The adaptive dynamics of such a game is therefore determined by the differential equation

$$\dot{x} = [\psi(x,x) - \phi(x,x)]\theta'_k(0) + \frac{1}{2} \frac{\partial}{\partial y} [\phi(y,x) + \psi(y,x)] \Big|_{y=x}. \quad (17)$$

The existence of singular strategies x^* in games of this form, and the particular characteristics of any such x^* , depend on $f_x(y)$ and $D(x)$, and thus ultimately on the specific functions $\phi(x,y)$ and $\psi(x,y)$. We shall now apply these results to the TD and MEC games.

Adaptive Dynamics of the Traveler's Dilemma Game

Let us first analyze the TD game with the payoff function given by (11). Consider a monomorphic population of x strategists, i.e., a population in which every individual claims amount x . It follows from (2) and (11) that for the TD game $\phi(x,y) = x + R$ and $\psi(x,y) = y - R$. Thus, the selection gradient $D(x)$ is given by

$$D(x) = \frac{1}{2} - 2R\theta'(0), \quad (18)$$

and the adaptive dynamics of x is consequently determined by

$$\dot{x} = \frac{1}{2} - 2R\theta'(0). \tag{19}$$

Since $D(x)$ does not depend on x , there are no singular strategies, and thus there is no possibility of exotic evolutionary outcomes, such as evolutionary branching. The evolutionary dynamics of an initial strategy x is determined by the sign of $D(x)$.

If $D(x) < 0$ (i.e., if $R > \frac{1}{4\theta'_k(0)}$), then x will evolve to 0. If on the other hand, $D(x) > 0$ (i.e., $R < \frac{1}{4\theta'_k(0)}$) then x will evolve to 1.

Thus, this adaptive dynamics analysis implies that the players of the smoothed TD game will evolve to make low claims if $R > \frac{1}{4\theta'_k(0)}$, and, conversely, evolve to make high claims if $R < \frac{1}{4\theta'_k(0)}$. We note that for the 1-parameter family of smoothing functions defined by (9), this criterion takes the following form: claims will evolve to low levels if $R > \frac{1}{k}$, and evolve to high levels if $R < \frac{1}{k}$.

Adaptive Dynamics of the Minimum Effort Coordination Game

We next turn to the MEC game where the payoff function is given by (14). Consider a monomorphic population of x strategists, i.e., a population in which every individual puts in x amount of effort. It follows from (2) and (13) that for the MEC game $\phi(x,y) = x - cx$ and $\psi(x,y) = y - cx$. Thus, the selection gradient $D(x)$ is given by

$$D(x) = \frac{1}{2} - c, \tag{20}$$

and the adaptive dynamics of x is therefore determined by

$$\dot{x} = \frac{1}{2} - c. \tag{21}$$

Again, since $D(x)$ does not depend on x , there are no singular strategies. Also, rather remarkably, the adaptive dynamics of x is independent of the smoothing function θ_k . The evolution of an initial strategy x is once again determined by the sign of $D(x)$. If $D(x) < 0$ (i.e., if $c > \frac{1}{2}$) then x will evolve to 0, and if $D(x) > 0$ (i.e., if $c < \frac{1}{2}$) then x will evolve to 1. Therefore, this adaptive dynamics analysis implies that the players' strategies in the smoothed MEC game will evolve to low efforts if the effort cost c is greater than $\frac{1}{2}$, and to high efforts if $c < \frac{1}{2}$. Since the adaptive dynamics of the smoothed MEC game is independent of the smoothing function, these results hold for any smoothing of the game.

We note that the behavior predicted by adaptive dynamics for the smoothed TD and MEC games is in accord with that observed for the TD and MEC games in experiments.

Agent-Based Simulations

In this section, we define a stochastic agent-based model which allows the evolutionary dynamics of the TD and MEC games to be

studied both for random interactions between members of the population and for more complex interaction patterns in the population. The evolutionary dynamics of simple social dilemmas, such as the Prisoner's Dilemma and the Snowdrift game, have been well-studied for populations with a variety of complex interaction patterns [42–50].

Consider a population consisting of n individuals, labeled $i = 1, \dots, n$. Since we wish to allow the possibility of complex population structures, we identify the population with the set of vertices in a graph Γ . The structure of Γ determines which individuals in the population can interact. Strictly speaking, two graphs are required to specify the evolutionary dynamics: an interaction graph, Γ_I , specifies that two individuals in the population can interact by playing the game only if they are adjacent in Γ_I , and an updating graph, Γ_U , specifies that an individual in the population can update its strategy by comparing its state to the states only of those individuals adjacent to it in Γ_U . Here, for simplicity, we shall assume that the interaction and updating graphs are the same i.e., $\Gamma_I = \Gamma_U = \Gamma$. Given an individual $i \in \Gamma$, the set of neighbors of i (i.e., the set of individuals adjacent to i in Γ) will be denoted by $N(i)$.

The agent-based model is defined as a stochastic process on Γ . Let us fix either the TD or the MEC game as the game under consideration. We begin with a monomorphic population, i.e., each individual in the population starts out with the same initial strategy randomly picked from a uniform distribution. At each time step t , we carry out a round of asynchronous interactions followed by a round of asynchronous updates. Each of these rounds involves sampling the population with replacement.

During an interaction step, we randomly pick an individual $i \in \Gamma$ and an individual $j \in N(i)$, and let the two individuals play the game against each other. If x and y denote the strategies of i and j , respectively, then the payoff P_i received by the focal individual i is given by either equation (11) or (14), depending on the game under consideration. This procedure is repeated n times.

During an update step, we randomly pick an individual $i \in \Gamma$ and an individual $j \in N(i)$. If P_i and P_j denote the payoffs of i and j , respectively, then with probability $p_{i \leftarrow j}$ given by

$$p_{i \leftarrow j} = \frac{1}{1 + e^{-\beta(P_j - P_i)}}, \tag{22}$$

the focal individual i will inherit j 's strategy. This update rule is often referred to as the Fermi rule. The parameter $\beta > 0$ is the "selection strength" of the update rule. The update procedure is repeated n times.

Mutations are incorporated in the update procedure in the following way: when according to the update rule (22) i 's strategy would be replaced by j 's, then with probability μ , i 's strategy is instead replaced by a strategy picked randomly from a normal distribution with mean equal to j 's strategy and standard deviation σ . Carrying out n interaction steps followed by n update steps constitutes a single generation of the evolutionary dynamics.

We note here that the results of our agent-based simulations (described in detail in the next section) are robust to variations in the update rule. For example, in addition to employing the Fermi update rule (22), we have also simulated the agent-based model using the replicator update rule, in which the probability $p_{i \leftarrow j}$ that the focal individual $i \in \Gamma$ inherits individual j 's strategy (with $j \in N(i)$) is given by

$$P_{i \leftarrow j} = \begin{cases} 0 & \text{if } P_i \geq P_j \\ \frac{\beta(P_j - P_i)}{M - m} & \text{otherwise,} \end{cases} \quad (23)$$

where $M = \max_{k \in \Gamma} P_k$, and $m = \min_{k \in \Gamma} P_k$. We find that the evolutionary dynamics of the smoothed TD and MEC games is the same irrespective of which of these update rules we employ. The results presented in the next section on the evolutionary dynamics of the TD and MEC games arise from simulations using the Fermi update rule (22).

Results

In this section we present the results of agent-based simulations for the TD and MEC games. For both the smoothed TD and MEC games the agent-based model described in the previous section was simulated (using the Fermi update rule (22) and the smoothing function (9)) on the following graphs (see, for example, [51–54].): a complete graph (which models a randomly-interacting population); a random regular graph of degree 10; a scale-free graph with mean degree 10; and two-dimensional lattice graphs with 4 and 8 neighbors, respectively, and periodic boundary conditions. The games were simulated for 20000 generations. The parameter values used for the simulations were: population size, $n = 1024$ for the lattice graphs and $n = 1000$ for the other graphs; mutation rate, $\mu = 0.05$; standard deviation for mutations, $\sigma = 0.005$; smoothing parameter, $k = 1.5, 2, 3$ for the TD game and $k = 0.5, 1, 2$ for the MEC game; and selection strength, $\beta = 5$.

Traveler's Dilemma Game

Figure 2 shows the variation of the average claims \bar{x} made by individuals over the last 20% of 20000 generations with the reward/punishment parameter R , for different values of the smoothing parameter k , on the following graphs: (a) a complete graph, (c) a random regular graph, (d) a scale-free graph, (e) a 2D lattice graph with 4 neighbors, and (f) a 2D lattice graph with 8 neighbors. The R value was varied from 0 to 1 in steps of 0.05, and each data point was obtained from an average of 10 runs of the model. It is apparent from these results that, for each value of k , the claims are high when $R < \frac{1}{k}$ and low when $R > \frac{1}{k}$, exactly as predicted by the adaptive dynamics analysis. This behavior is in good qualitative agreement with the results obtained for the TD game in experiments. Furthermore, these simulation results suggest that network structure has very little effect on the evolutionary dynamics of the game.

Figure 2(b) shows the variation in the number of individuals with the claims x they make, when $R = 0.15$ and $R = 0.85$, on a complete graph, with $k = 2$. Individuals make higher claims when $R = 0.15$, as indicated by the blue bars, and make lower claims when $R = 0.85$, as indicated by the green bars. This result is in good agreement with the empirical results shown in Figure 1(a).

We have also simulated the game using the discontinuous form of the payoff function (10), and the results are shown in blue (labeled $k = \infty$) in panels (a), (c), (d), (e), and (f) of Figure 2. In this case, the claims are consistently low for all values of R . This result, which is consistent with the prediction of the adaptive dynamics analysis of the smoothed TD game in the limit $k \rightarrow \infty$, is in agreement with the prediction of classical game theory for the original TD game. Thus, for the TD game studied here, the smoothing of the payoff function is necessary to explain the empirically observed behavior.

Minimum Effort Coordination Game

Figure 3 shows the variation of the average effort levels \bar{x} of the individuals over the last 20% of 20000 generations with the effort cost parameter c , for various values of the smoothing parameter k , on the following graphs: (a) a complete graph, (c) a random regular graph, (d) a scale-free graph, (e) a 2D lattice graph with 4 neighbors, and (f) a 2D lattice graph with 8 neighbors. The c value was varied from 0 to 1 in steps of 0.05, and each data point was obtained from an average over 10 runs of the model. It is clear from these results that, for each value of c , the effort levels are high when $c < \frac{1}{2}$ and low when $c > \frac{1}{2}$, exactly as suggested by the adaptive dynamics analysis. These results are in good qualitative agreement with the empirically observed behavior in experiments. The results of these agent-based simulations clearly show the independence of the evolutionary dynamics of the smoothed MEC game on the smoothing function (i.e., the independence on k), which is predicted by the adaptive dynamics analysis. These results also suggest that the effect of network structure on the dynamics of the smoothed MEC game is negligible.

Figure 3(b) shows the variation in the number of individuals with their effort levels x , when $c = 0.15$ and $c = 0.85$, on a complete network, with $k = 2$. Individuals expend more effort when $c = 0.15$, as indicated by the blue bars, and expend lower effort when $c = 0.85$, as indicated by the green bars. This result is in good agreement with the empirical results shown in Figure 1(b).

We also simulated the model using the discontinuous form of the payoff function (13), and the result is shown in blue ($k = \infty$) in plots (a), (c), (d), (e), and (f) of Figure 3. In this case, there is still a transition from high to low effort levels as the effort cost increases, however, it is typically not as sharp as for the smoothed game (with any value of k). An intriguing exception to this pattern is found in the case of two-dimensional lattice graphs, for which the evolutionary dynamics of the game with discontinuous payoff function is essentially identical to that of the smoothed game (for any value of the smoothing parameter). We tentatively conjecture that the feature of lattice graphs which is responsible for this effect is that they possess non-trivial clustering coefficients ([51–54]) — in contrast to the other graphs we have considered, which have zero clustering. The potential effect of the clustering coefficient on the evolutionary dynamics of the MEC game defined on graphs appears to be an interesting topic for future research.

Discussion

In this work we have proposed simple and natural continuous-strategy versions of the classical discrete-strategy TD and MEC games. We have modeled these games as continuous games with smooth payoff functions, where the smoothing accounts for the effects of errors in the perception and/or implementation of individuals actions. The smoothed TD and MEC games can be effectively analyzed using adaptive dynamics, which shows that the predicted evolutionary dynamics of these games is in accord with the behavior observed in empirical studies of the TD and MEC games. In addition, we have studied the evolutionary dynamics of the smoothed TD and MEC games using agent-based simulations. These simulations have been performed both for populations of randomly-interacting agents and for populations with more complex interaction patterns, represented by graphs of varying topologies. These simulation results are in agreement both with the analytical adaptive dynamics results, and also with the experimentally observed behavior.

For the smoothed TD game, we find both from the adaptive dynamics analysis and from the agent-based simulations, that

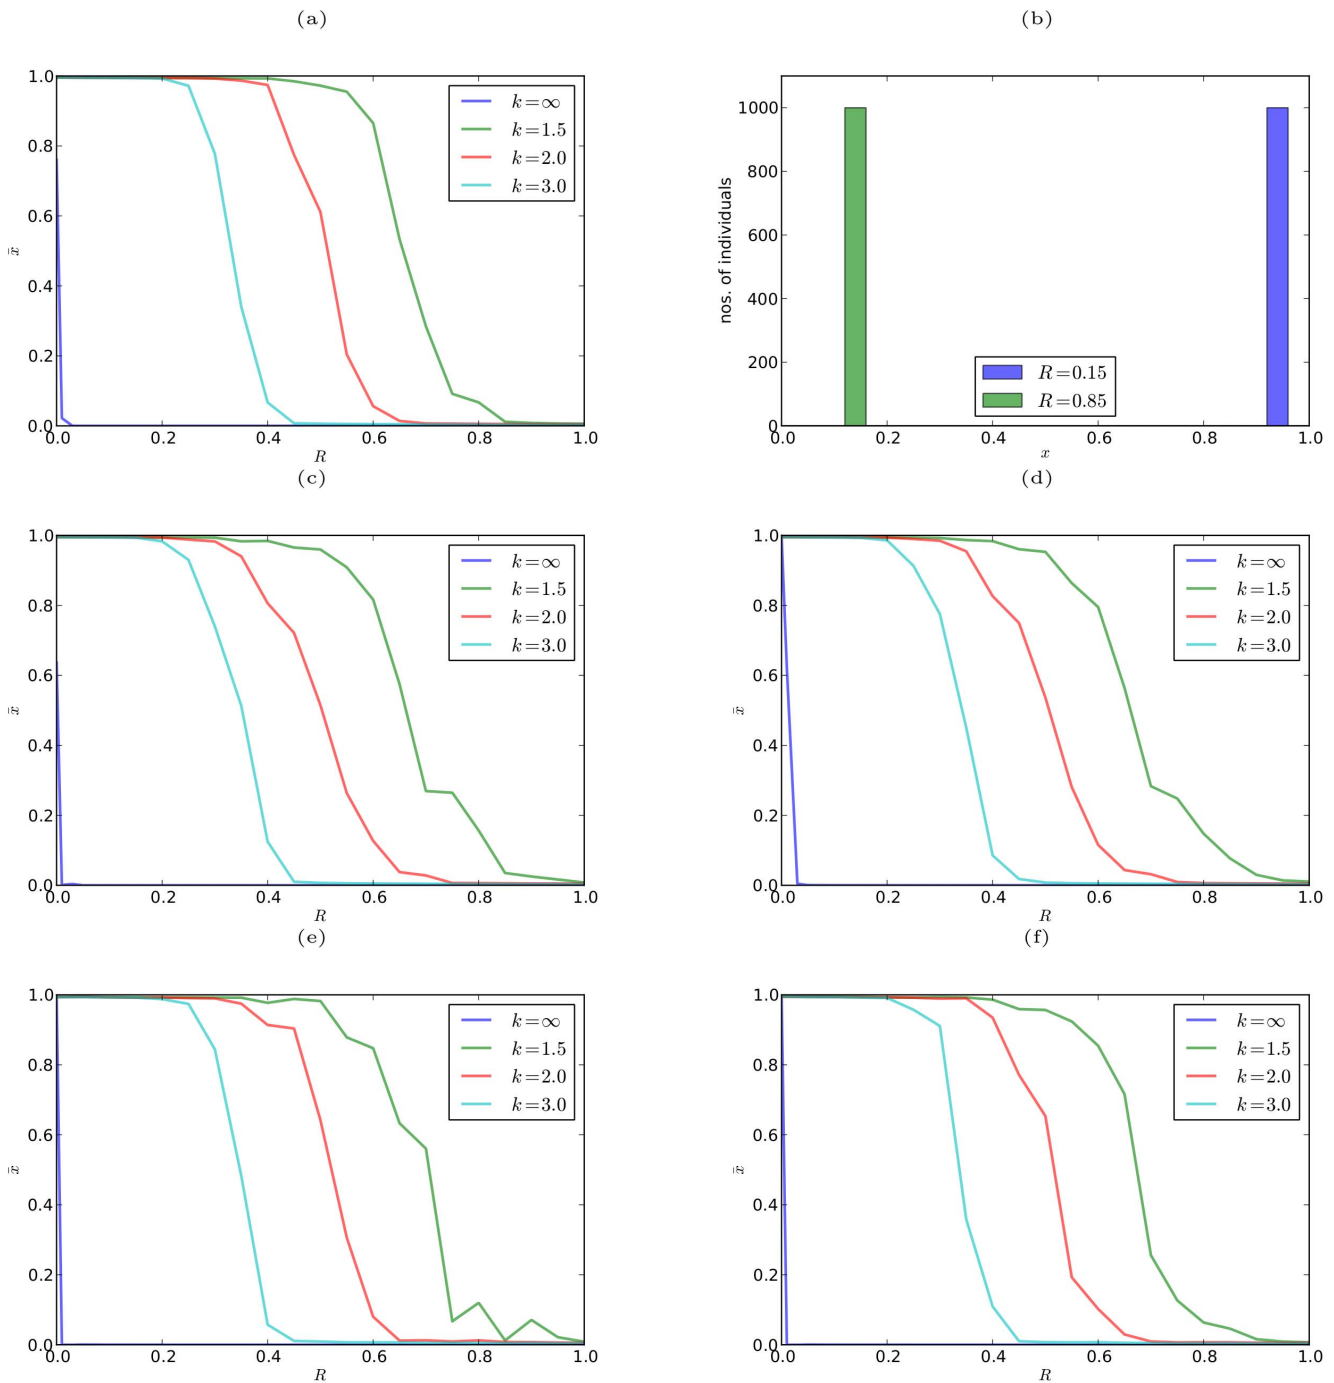


Figure 2. Results from simulating the TD game. (a)(c)(d)(e)(f) Average claims \bar{x} in the smoothed TD game over the last 20% of 20000 generations versus the reward/punishment parameter R for different values of the smoothing parameter k , on a complete graph (a), a random regular graph with degree 10 (c), a scale-free graph with mean degree 10 (d), 2D lattice graph with 4 neighbors (e), and 2D lattice graph with 8 neighbors (f). Parameter values: $n = 1024$ for lattice networks and $n = 1000$ for other networks, $\mu = 0.05$, $\sigma = 0.005$, and $\beta = 5$. (b) Number of individuals versus their claims x , when $R = 0.15$ and $R = 0.85$, on a complete graph with parameter values: $n = 1000$, $\mu = 0.05$, $\sigma = 0.005$, $k = 2$, and $\beta = 5$. doi:10.1371/journal.pone.0093988.g002

claims vary with the reward/punishment parameter R in a fashion that is in excellent agreement with the empirically observed behavior: low values of R result in high claims and high values of R result in low claims. We recover the classical game theory result that claims in the TD game with discontinuous payoff remain low for all values of R by considering the limit in which the smoothing parameter $k \rightarrow \infty$. Different interaction patterns among the

individuals playing the smoothed TD game, as represented by studying the game on graphs of different topologies, appears to have little effect on the evolutionary dynamics of the game.

We find similarly satisfactory results for the smoothed MEC game. The analysis, both analytical and through simulations, again yields results in good agreement with experiment: high effort levels are found for low effort cost c and low effort levels occur for high

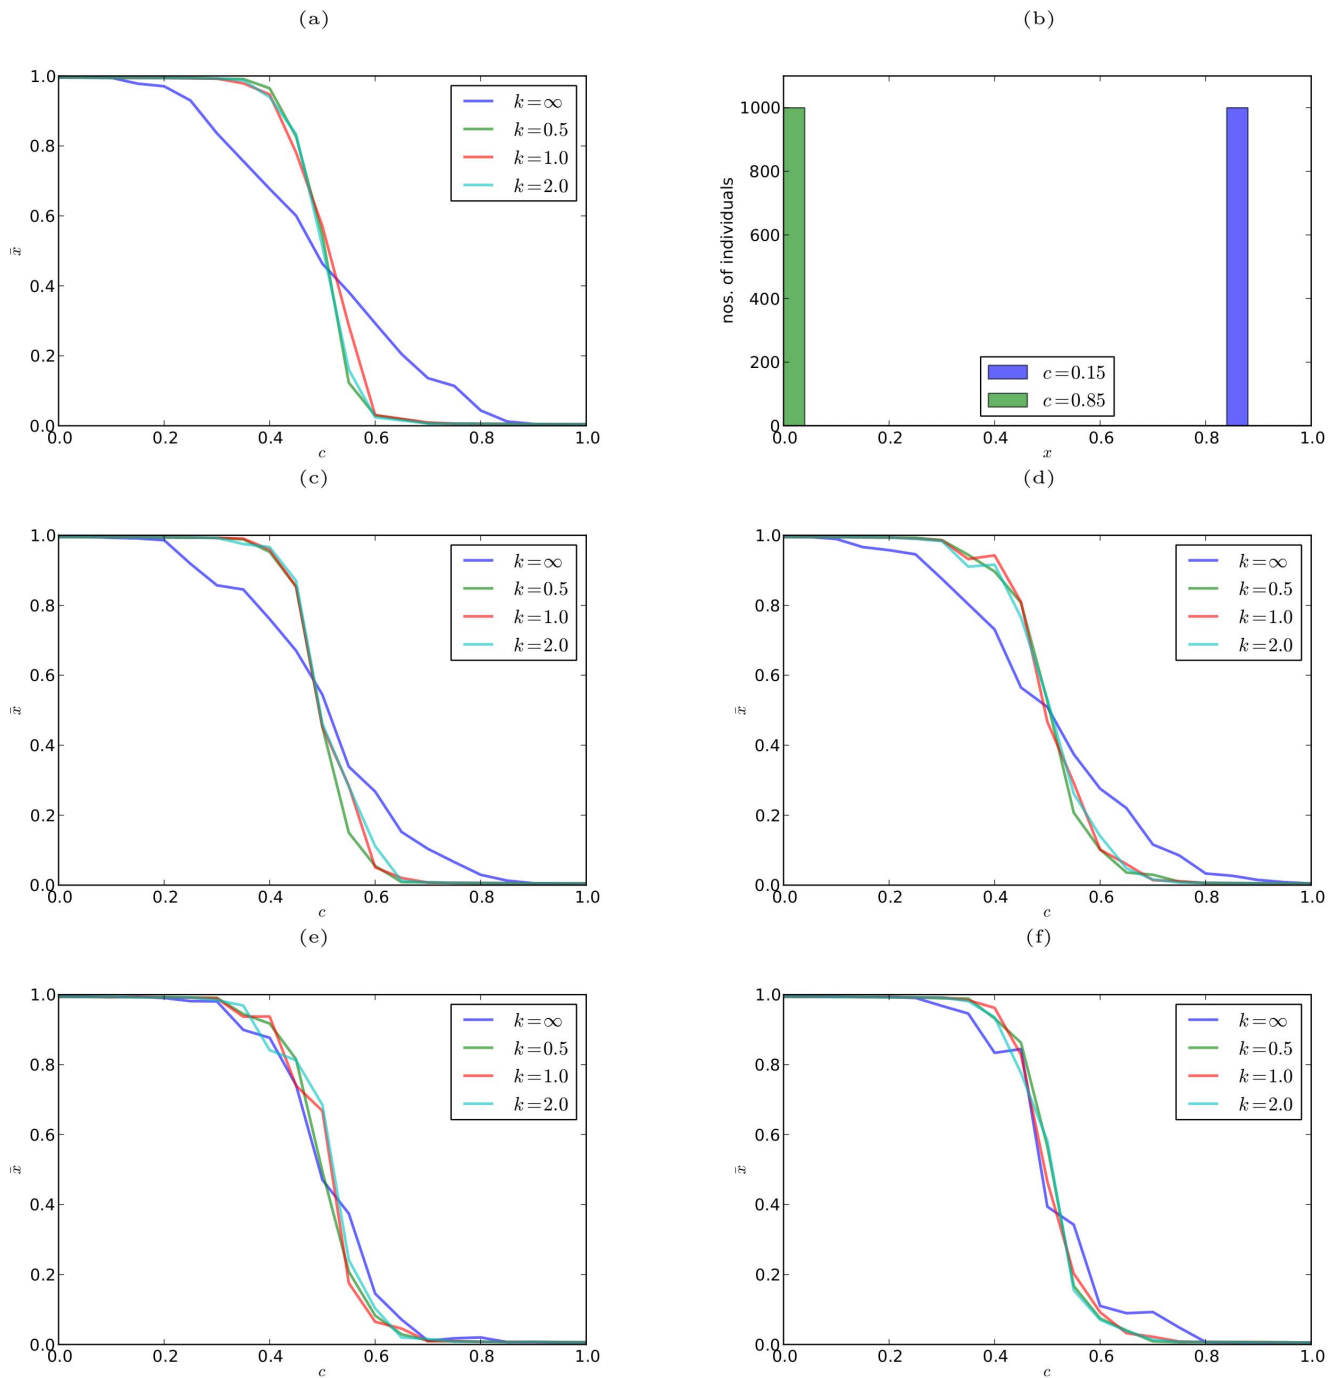


Figure 3. Results from simulating the MEC game. (a)(c)(d)(e)(f) Average effort levels \bar{x} in the smoothed MEC game over the last 20% of 20000 generations versus the effort cost parameter c for different values of the smoothing parameter k , on a complete graph (a), a random regular graph with degree 10 (c), a scale-free graph with mean degree 10 (d), 2D lattice graph with 4 neighbors (e), and 2D lattice graph with 8 neighbors (f). Parameter values: $n = 1024$ for lattice graphs and $n = 1000$ for other graphs, $\mu = 0.05$, $\sigma = 0.005$, and $\beta = 5$. (b) Number of individuals versus their effort levels x , when $c = 0.15$ and $c = 0.85$, on a complete graph with parameter values: $n = 1000$, $\mu = 0.05$, $\sigma = 0.005$, $k = 2$, and $\beta = 5$. doi:10.1371/journal.pone.0093988.g003

effort cost. We again find that when the smoothed MEC game is formulated on graphs of differing topology the topological type has no significant effect on the evolutionary dynamics of the game.

The methods introduced in this paper are quite general and can be applied to a wide variety of continuous-strategy games with discontinuous payoff functions. Important examples of other games that can be fruitfully studied using these methods include

the Bertrand Duopoly model [38] and the War of Attrition game [6]. Here we will only briefly discuss the application of our methods to these two games — detailed accounts will be given elsewhere.

In the classical Bertrand Duopoly (BD) model [38] one considers the interactions between two firms that produce a homogeneous product. The strategy of each firm is the unit price

that they set for their product. It is assumed that customers will buy a quantity $d(p)$ (where $d(p)$ is the demand function) from the firm with the lower price p , and will buy nothing from the firm with the higher price. If both firms set the same price then it is assumed that demand is split equally between them. It is also assumed that both firms have the same marginal cost c . Thus, the payoff obtained by one firm in its interaction with the other is a discontinuous function of the difference between the prices set by the two firms. This discontinuous payoff function can be smoothed in exactly the manner described in this paper to yield a smoothed game. In this case the smoothing function represents the probability that customers buy the product from the firm with the lower price as opposed to the firm with the higher price. In the limit that this probability tends to one, the smoothed game approaches the classical game.

The smoothed BD model can be analyzed using adaptive dynamics just as we have done here for the TD and MEC games. It may be shown for an arbitrary smoothing, that given a linear demand function $d(p)$ there exists a unique singular strategy p^* (in the domain of interest), that is strictly greater than the marginal cost c . Furthermore, it may also be shown that the singular strategy p^* is always both convergent stable and an ESS. Thus, the evolutionary dynamics of the smoothed BD game results in the prices set by both firms converging to the level p^* , which is strictly greater than the marginal cost. At the evolutionary equilibrium p^* both firms obtain a positive payoff. Thus, the introduction of the smoothed BD model effectively resolves the Bertrand Paradox that occurs in the classical model. Moreover, our method extends to the case of an arbitrary number of firms and sheds light on the results obtained in [55].

References

- Kollock P (1998) Social dilemmas: The anatomy of cooperation. *Annual Review of Sociology* 24: 183–214.
- Axelrod R (2006) *The Evolution of Cooperation*. Basic books.
- Hardin G (1968) The tragedy of the commons. *Science* 162: 1243–1248.
- Flood MM (1958) Some experimental games. *Management Science* 5: 5–26.
- Hamilton W, Axelrod R (1981) The evolution of cooperation. *Science* 211: 1390–1396.
- Maynard Smith J (1982) *Evolution and the Theory of Games*. Cambridge University Press.
- Skyrms B (2003) *The Stag Hunt and the Evolution of Social Structure*. Cambridge University Press.
- Isaac R, Walker J (1988) Group size effects in public goods provision: The voluntary contributions mechanism. *The Quarterly Journal of Economics* 103: 179–199.
- Fehr E, Gächter S (2002) Altruistic punishment in humans. *Nature* 415: 137–140.
- Killingback T, Doebeli M, Knowlton N (1999) Variable investment, the continuous prisoner's dilemma, and the origin of cooperation. *Proceedings of the Royal Society of London Series B: Biological Sciences* 266: 1723–1728.
- Doebeli M, Hauert C, Killingback T (2004) The evolutionary origin of cooperators and defectors. *Science* 306: 859–862.
- Killingback T, Doebeli M, Hauert C (2010) Cooperation and defection in the tragedy of the commons. *Biological Theory* 5: 3–6.
- Killingback T, Bieri J, Flatt T (2006) Evolution in group-structured populations can resolve the tragedy of the commons. *Proceedings of the Royal Society of London Series B: Biological Sciences* 273: 1477–1481.
- Holt C, Roth A (2004) The Nash equilibrium: A perspective. *Proceedings of the National Academy of Sciences* 101: 3999–4002.
- Basu K (1994) The traveler's dilemma: Paradoxes of rationality in game theory. *The American Economic Review* 84: 391–395.
- Basu K (2007) The traveler's dilemma. *Scientific American Magazine* 296: 90–95.
- Goeree J, Holt C (2001) Ten little treasures of game theory and ten intuitive contradictions. *The American Economic Review* 91: 1402–1422.
- Van Huyck J, Battalio R, Beil R (1990) Tacit coordination games, strategic uncertainty, and coordination failure. *The American Economic Review* 80: 234–248.
- Capra C, Goeree J, Gomez R, Holt C (1999) Anomalous behavior in a traveler's dilemma? *The American Economic Review* 89: 678–690.
- Goeree J, Holt C (1999) Stochastic game theory: For playing games, not just for doing theory. *Proceedings of the National Academy of Sciences* 96: 10564–10567.
- Becker T, Carter M, Naeve J (2005) Experts Playing the Traveler's Dilemma. *Inst. für Volkswirtschaftslehre, Univ.*
- Basu K, Becchetti L, Stanca L (2011) Experiments with the Travelers Dilemma: welfare, strategic choice and implicit collusion. *Social Choice and Welfare* 37: 575–595.
- Hofbauer J, Sigmund K (1998) *Evolutionary games and population dynamics*. Cambridge University Press.
- Manapat M, Rand D, Pawlowitsch C, Nowak M (2012) Stochastic evolutionary dynamics resolve the Traveler's Dilemma. *Journal of Theoretical Biology* 303: 119–127.
- Riehl C (2011) Living with strangers: direct benefits favour non-kin cooperation in a communally nesting bird. *Proceedings of the Royal Society B: Biological Sciences* 278: 1728–1735.
- Riehl C, Jara L (2009) Natural history and reproductive biology of the communally breeding Greater (Ani Crotophaga) major at Gatún Lake, Panama. *The Wilson Journal of Ornithology* 121: 679–687.
- Bryant J (1983) A simple rational expectations Keynes-type model. *The Quarterly Journal of Economics* 98: 525–528.
- Anderson S, Goeree J, Holt C (2001) Minimum-effort coordination games: Stochastic potential and logit equilibrium. *Games and Economic Behavior* 34: 177–199.
- Halpern J, Pass R (2012) Iterated regret minimization: A new solution concept. *Games and Economic Behavior* 74: 184–207.
- Capraro V (2013) A model of human cooperation in social dilemmas. *PLoS ONE* 8: 72427.
- Capraro V, Venanzi M, Polukarov M, Jennings N (2013) Cooperative equilibria in iterated social dilemmas. *Proceedings of the 6th International Symposium in Algorithmic Game Theory* 8146: 146–158.
- Li R, Yu J, Lin J (2013) Evolution of cooperation in spatial Traveler's Dilemma game. *PLoS ONE* 8: 58597.
- Camerer C, Ho T, Chong J (2004) A cognitive hierarchy model of games. *The Quarterly Journal of Economics* 119: 861–898.
- Costa-Gomes M, Crawford V, Broseta B (2001) Cognition and Behavior in Normal-Form Games: An Experimental Study. *Econometrica* 69: 1193–1235.
- McKelvey R, Palfrey T (1995) Quantal response equilibria for normal form games. *Games and Economic Behavior* 10: 6–38.
- Stahl D, Wilson P (1994) Experimental evidence on players' models of other players. *Journal of Economic Behavior and Organization* 25: 309–327.

37. Wright J, Leyton-Brown K (2010) Beyond Equilibrium: Predicting Human Behavior in Normal- Form Games. Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence: 901–907.
38. Tirole J (1988) The Theory of Industrial Organization. MIT press.
39. Metz J, Geritz S, Meszéna G, Jacobs F, Van Heerwaarden J (1996) Adaptive dynamics, a geometrical study of the consequences of nearly faithful reproduction. *Stochastic and Spatial Structures of Dynamical Systems* 45: 183–231.
40. Geritz S, Metz J, Kisdi É, Meszéna G (1997) Dynamics of adaptation and evolutionary branching. *Physical Review Letters* 78: 2024–2027.
41. Geritz S, Meszéna G, Metz J (1997) Evolutionarily singular strategies and the adaptive growth and branching of the evolutionary tree. *Evolutionary Ecology* 12: 35–57.
42. Nowak M, May R (1992) Evolutionary games and spatial chaos. *Nature* 359: 826–829.
43. Killingback T, Doebeli M (1996) Spatial evolutionary game theory: Hawks and Doves revisited. Proceedings of the Royal Society of London Series B: Biological Sciences 263: 1135–1144.
44. Nakamaru M, Matsuda H, Iwasa Y (1997) The evolution of cooperation in a lattice-structured population. *Journal of Theoretical Biology* 184: 65–81.
45. Van Baalen M, Rand D (1998) The unit of selection in viscous populations and the evolution of altruism. *Journal of Theoretical Biology* 193: 631–648.
46. Ifiti M, Killingback T, Doebeli M (2004) Effects of neighbourhood size and connectivity on spatial Continuous Prisoner's Dilemma. *Journal of Theoretical Biology* 231: 97–106.
47. Hauert C, Doebeli M (2004) Spatial structure often inhibits the evolution of cooperation in the snowdrift game. *Nature* 428: 643–646.
48. Santos F, Pacheco J (2005) Scale-free networks provide a unifying framework for the emergence of cooperation. *Physical Review Letters* 95: 98–104.
49. Ohtsuki H, Hauert C, Lieberman E, Nowak M (2006) A simple rule for the evolution of cooperation on graphs and social networks. *Nature* 441: 502–505.
50. Szabo G, Fáth G (2007) Evolutionary games on graphs. *Physics Reports* 446: 97–216.
51. Watts D, Strogatz S (1998) Collective dynamics of 'small-world' networks. *Nature* 393: 440–442.
52. Barabási A, Albert R (1999) Emergence of scaling in random networks. *Science* 286: 509–512.
53. Albert R, Barabási A (2002) Statistical mechanics of complex networks. *Reviews of Modern Physics* 74: 47–97.
54. Newman M (2010) *Networks: An Introduction*. Oxford University Press.
55. Dufwenberg M, Gneezy U (2000) Price competition and market concentration: an experimental study. *International Journal of Industrial Organization* 18: 7–22.