Heinrich H. Nax, Ryan O. Murphy, and Dirk Helbing

Nash Dynamics, Meritocratic Matching, and Cooperation

Abstract: John F Nash (1950) proposed dynamics for repeated interactions according to which agents myopically play individual best-responses against their observations of other agents' past play. Such dynamics converge to Nash equilibria. Without suitable mechanisms, this means that best-response dynamics can lead to low levels of cooperative behavior and thus to inefficient outcomes in social dilemma games. Here, we discuss the theoretical predictions of these dynamics in a variety of social dilemmas and assess these in light of behavioral evidence. We particularly focus on "meritocratic matching", a class of mechanisms that leads to both low cooperation (inefficient) and high cooperation (near-efficient) equilibria (Gunnthorsdottir et al. 2010; Nax, Murphy, and Helbing 2014; Nax et al. 2015). Most behavioral theories derived from related social dilemmas cannot explain the behavioral evidence for this class of games, but Nash dynamics provide a very satisfactory explanation. We also argue that Nash dynamics provide a parsimonious account of behavioral results for several different social dilemmas, with the exception of the linear public goods game.

1 Introduction

Without an appropriate institution, norm or mechanism, cooperation in repeated social dilemma interactions tends to deteriorate over time. Evidence of such phenomena is not restricted to controlled laboratory experiments (see, for example, Ledyard 1995 and Chaudhuri 2011 for reviews); it is also a widely recorded phenomenon in real-world interactions such as collective action and common-pool resource management (Olson 1965; Ostrom 1990; Ostrom 2000; Ostrom 2005; Ostrom, Walker, and Gardner 1994; Ostrom and Walker 2013). The long-term fate of these collective interactions is often a tragedy of the commons (Hardin 1968), a situation in which narrow self-interest and rational behavior lead to collectively worse outcomes than could be achieved by collective action.

Game theory (von Neumann and Morgenstern 1944), especially non-cooperative game theory (Nash 1950; Nash 1951), has made formal analyses of strategic interactions possible. This formal framework allows the identification of the key characteristics that determine whether interactions will lead to efficient outcomes or succumb to the tragedy of the commons.

Large parts of the game-theoretic inquiry are concerned with static models. In this chapter, we focus on dynamic models, particularly best-response dynamics as first introduced in Nash's Ph.D. dissertation (Nash 1950). According to such dynamics, indi-

https://doi.org/10.1515/9783110472974-021

viduals play a repeated game and adapt by myopically best responding, or choosing the strategy that maximizes their current-period payoff in response to the strategies they observed other agents to adopt in the past. Such dynamics converge to a Nash equilibrium, as articulated in Nash's dissertation. This part of Nash's dissertation is not as well-known because the short section of the dissertation that contained this material was omitted in the publication (Nash 1951; see Young 2011 for a discussion).

More broadly, Nash dynamics are part of the strand of evolutionary game theory models that are not strictly based on the classical replicator argument. Other important dynamics that are closely related to Nash dynamics include fictitious play, stochastic fictitious play, and other forms of belief-based learning models. It is worth emphasizing again that all these dynamic approaches are about noisily best-responding agents, usually presuming agents are driven by narrow self-interest, and that, even holding preferences constant, we can go far in understanding and modeling emergent behavior when considering such dynamics.

The dynamic modeling of emergent behavior is a fruitful approach in two ways. First, it is often quite powerful: we can explain why certain outcomes are selected rather than others by looking at the long-term outcomes of the noisy dynamics. Second, dynamic modeling is general: we do not need a different preference theory for each different game, or a preference model at the individual level for all these different games. This is not to argue that individual level preference models are wrong or useless, but they may not be a good first step in trying to understand complicated social interactions and associated emergence behavior. The parsimony of the dynamic model makes it a promising first step and, as we show, can be applied broadly. In the next section, we apply a Nash dynamics analysis to four types of social dilemma game, for each of which there are many real-world examples.

(1) *The linear public goods game.*
 Individuals decide separately whether, and if so how much, to contribute to the provisioning of a public good. Each individual contribution leads to a greater collective benefit, but the individual cost of its provision exceeds the private benefit. There is therefore a clash between private and collective interests. Not contributing is the dominant individual strategy, but full contribution by everyone maximizes collective payoffs.
 This game, with a simple linear payoff structure, was first studied as the voluntary contribution mechanism (VCM) applied to the linear public goods game (L-PGG) by Marwell and Ames (1979) (see also Issac, McCue, and Plott 1985). The unique Nash equilibrium of the VCM in the context of L-PGG is universal freeriding/non-contribution, while the outcome that maximizes collective payoffs is universal full-contribution.

(2) *The step-level public good.*
 In this game, a certain minimum number (or a certain level) of contributions is needed to reach a threshold at or above which the public good is provided. Reaching the threshold requires contributions by two or more players. An individual is

best off if the public good is provided but he himself freerides. The second-best outcome is for the public good to be provided and for the individual to have also contributed himself. Third-best outcome is for the public good not to be provided, without the individual contributing himself. The worst outcome for the individual is to contribute but for the total contributions to fall short of the threshold, leading to the public good not being provided.

Palfrey and Rosenthal (1984) introduced this class of game, which has since been studied experimentally (e.g., Kragt, Orbell, and Dawes 1983; Rapoport and Suleiman 1993). There are two types of pure Nash equilibria in this game: one where the threshold is exactly met, which is also the outcome that potentially maximizes collective payoffs, and one where no one contributes, that is, non-provision of the public good. One therefore has a coordination problem: who, if anyone, will contribute?

(3) *The volunteer's dilemma.*

Exactly one volunteer is needed to provide the public good. The individual is best-off if the public good is provided by someone else (i.e., someone else volunteers). However, in case no one else volunteers, the individual best response is to do it (i.e., to be the volunteer).

Diekmann (1985) introduced this game. The pure Nash equilibria in this game are asymmetric, with any one of the players volunteering. There is also a mixed equilibrium that results in positive probabilities of volunteering. The socially desirable outcome, in terms of total collective payoffs, is for the volunteer to be the one who has the lowest cost of volunteering, but the problem is again one of coordination: who will volunteer?

(4) *Group-based meritocratic matching.*

Individuals jointly create a good, which is of a public-goods character, within several separate "clubs" (as in Buchanan 1965): that is, the benefits from the public goods provided in each club do not transcend club boundaries. Inside the club, the same structure as in L-PGG prevails, but admittance to clubs is based on contribution decisions. We refer to this contribution-based group admittance as "meritocratic matching".

Gunnthorsdottir et al. (2010), Nax, Murphy, and Helbing (2014), and Nax et al. (2015) formulated such "meritocratic matching" mechanisms theoretically and empirically. As in the step-level public goods game, the game potentially (whether this is the case depends on various parameters of the game) has two types of pure Nash equilibria: one (which often but not always exists) with many contributions, which is near-efficient, and the other (which always exists) with universal freeriding. Again, a coordination problem emerges and the crux of the strategic interaction is finding the cooperators/contributors.

Next, we will show that the kind of Nash dynamics discussed above provide a simple and parsimonious account of behavior in many social dilemma games. Only linear public goods games are an exception to this dynamics account. We believe this is notable because a large (perhaps disproportionate) amount of attention has been given to these games, which has led to a commonly held conclusion that more complex mechanisms like social preferences (or reciprocity, which is changes/dynamics in actions or social preferences) are required to understand cooperative behavior in social interactions. While this may be true for linear public goods games, this may not be so in many other games, as selfish but imperfect (noisy) agents interacting in particular settings lead to efficient outcomes, and yield various commonly observed dynamics.

The spontaneous invocation of social preferences in these other games is perhaps unnecessary, as other simpler mechanisms can account for much of the observed behavior from experiments. Moreover, these other social dilemma games contain multiple equilibria that better correspond to various interactions in the real world, as has been argued, for example, by various biologists studying the volunteer's dilemma (Raihani and Bshary 2011). Often, mechanisms transform interactions into coordination games from social dilemma games, and social dilemmas with multiple equilibria present agents again with a coordination problem. Take the example of cooperation emerging in the presence of punishment. If cooperation is upheld by punishment of defectors, it is an open question who will punish when defection occurs.

What is noteworthy is that Nash dynamics offer a solution to this coordination problem. The dynamics settle into equilibria whose "basins of attraction" are bigger than others. One example of this is the generalized Meritocracy games we developed. In these games, the near-efficient equilibria often have a large basin of attraction compared with the zero-contribution equilibrium, and this feature foments interesting dynamics. It establishes a "rational" pathway for narrowly self-interested agents to secure the tremendous efficiency gains available in this strategic context. Contrast this with L-PGG, where there is no such mechanism.

Taken all together, we see a variety of interesting social dilemmas in the wild that are not isomorphic to the standard PD game or the common L-PGG. We strongly encourage researchers to pay more attention to identifying, modeling, and studying these other social dilemmas that are different than the fruit flies we know well (e.g., PD and L-PGG). In addition, we particularly encourage the development of quantitative models of social dilemmas that facilitate the exploration of particular mechanisms (information, signaling, etc.) that can facilitate the producing of interesting (non-trivial) dynamics. Lastly, we encourage parsimonious approaches to understanding emergent dynamics. This implies starting with a concept like Nash dynamics (or stochastic fictitious play, or other simple myopic learning models) as a way to begin accounting for emergent behavior rather than evoking more complicated preference refinements. This is not to say that non-selfish preferences do not exist, or are not potentially important, but rather that there may be other simpler mechanisms that can account for the empirics, and Occam's razor directs us to look to these simpler mechanisms first.

The focus of this chapter will be on Nash dynamics and on *meritocratic matching*, which, broadly, falls under "assortative matching", a widespread phenomenon. In evolutionary biology, assortative matching is the key mechanism underlying various forms of kin selection (Hamilton 1964a; Hamilton 1964b); for example, via limited dispersal/locality (spatial interactions; Nowak and May 1992; Eshel, Samuelson, and Shaked 1998; Skyrms 2004) or greenbeard genes (Dawkins 1976; Fletcher and Doebeli 2009; Fletcher and Doebeli 2010). Similarly, assortative matching can be expressed via "homophily" (Alger and Weibull 2012; Alger and Weibull 2013; Xie, Cheng, and Zhou 2015).

More specifically, meritocratic matching is a mechanism that leads to assortativity of actions, rather than of genes or locations, and is particularly relevant for human interactions, where institutions exist that can determine who interacts with whom based on observable behavior. The "meritocratic" element is part of the institutional structure of the interaction, and not part of the decisions made by the involved individuals (even though one can think of it as an evolving institution: Nax and Rigos 2016). Examples include school/university admission, team-based payment schemes, and organizational selection and recruitment. These examples have in common that individuals gain access to better groups (leading to better payments) based on their own effort and performance. The crucial common characteristic here is that agents make a precommitted (irrevocable) investment decision that leads to assortative matching with potentially important payoff consequences at the end of the interaction (see Nax, Murphy, and Helbing 2014 for a more detailed discussion). In human interactions, meritocratic matching has been shown, theoretically and experimentally, to stabilize near-efficient contribution levels effectively (Gunnthorsdottir et al. 2010; Nax, Murphy, and Helbing 2014; Nax et al. 2015).

To further assist in understanding how meritocratic matching works, we shall explain briefly (before introducing the mechanism formally in a subsequent section) its basic principles. A population of agents is divided into several groups based on contribution decisions: contributors (freeriders) tend to be matched with other contributors (freeriders). Contribution decisions precede group matching, and meritocratic matching creates incentives to contribute to be matched with others doing likewise. The resulting structures of Nash equilibria are as follows. On the one hand, there exist multiple asymmetric near-efficient equilibria with many contributors and only a few freeriders. On the other hand, there continues to exist a Pareto deficient symmetric equilibrium in which all players freeride. It is noteworthy that the meritocracy mechanism enables high-efficiency equilibria to emerge even with narrowly self-regarding players. This stands in notable contrast to the vast majority of existing work on cooperation that invokes other-regarding preferences or reciprocity (i.e., contingent other-regarding preferences) as a means of achieving collective improvement.

To add to the discussion about such mechanisms, the purpose of this chapter is threefold. First, the consequences of "Nash dynamics" in various social dilemma games are reviewed and compared with experimental evidence. Second, these dynamics are used as the basis to explain, in an elemental way (retaining the assumption of self-regarding preferences), how the players coordinate playing strategies that yield a nearly perfectly efficient equilibrium under meritocratic matching. Finally, differences and similarities with L-PGG, step-level public goods, and volunteer's dilemmas are discussed.

The rest of this chapter is divided as follows. First, we introduce Nash dynamics. In Section 3, we introduce the classes of social dilemmas under consideration. Sections 4 to 5, respectively, present the predictions of the Nash dynamics, discuss existing experimental evidence, and provide alternative explanations. Section 6 concludes.

2 Nash dynamics

Evolutionary arguments form the backbone of much of the "emergence of cooperation" literature (Axelrod and Hamilton 1981; Axelrod 1984) related to social dilemmas, of which, for example, West, Griffin, and Gardner (2007), and West, Mouden, and Gardner (2011) provide reviews from an evolutionary biology perspective. Nash himself proposed a particular kind of dynamic justification for his equilibrium concept, and we will focus on these "Nash dynamics" in this chapter. Before we introduce them formally, however, we would like to provide a less formal overview.

First, let us recall the static justification of the *Nash equilibrium*: an outcome of a game is a Nash equilibrium if and only *if all strategies that are being played constitute mutual best replies to one another*. In other words, in an equilibrium it behooves no player to change their selected strategy unilaterally.

How is a Nash equilibrium reached?

One way to think about how a Nash equilibrium is arrived at, is as the outcome of an indefinitely repeated game. In that game, agents have the opportunity to revise their strategies over time in light of the past actions of others as played over iterations of the game. If all agents myopically choose to play best replies against their observations of other players' chosen strategies from the past, who themselves also play best replies given their observations, then such a process will lead to a Nash equilibrium. Certainly, any Nash equilibrium is an absorbing state of such a dynamic process as no player has an incentive to unilaterally choose another action given the other players' choices.

Which equilibria will emerge when there are multiple equilibria?

This equilibrium selection question has been addressed in the evolutionary games literature. One route of enquiry, not based on Nash dynamics, is concerned with *evolutionary stable strategies* (Taylor and Jonker 1978; Helbing 1992; Helbing 1996; Weibull

1995) based on the replication/imitation of strategies with higher *fitness* (Darwin 1871; Maynard Smith and Price 1973; Maynard Smith 1987). Such arguments are at the heart of evolutionary game theory (often abbreviated as EGT) as applied to biology.

In social scientific enquiries of human decision-making, beginning with the seminal contributions by Foster and Young (1990), Young (1993) and Kandori, Mailath, and Rob (1993), analyses take Nash dynamics as the baseline instead of replication/imitation which presumes only gradual adaptation. The crucial novelty is that "noise" is added in the sense of random behavioral deviations from the predominant best-response rule (Helbing 2010; Mäs and Helbing 2014). This added noise, in various forms, has generated sharp long-run predictions and led to a rich theoretical literature, where long-run predictions of which Nash equilibria will be selected depend crucially on the noise modeling (Bergin and Lipman 1996; Blume 2003). Recently, the underlying assumptions of noise are being investigated behaviorally using laboratory studies (Mäs and Nax 2015; Young 2015).

To date, most of this literature focuses on coordination games. In the context of social dilemma games, the connections between evolutionary explanations and behavioral/experimental studies are less direct: experimental evidence is explained by social preferences and social norms (e.g., Chaudhuri 2011), and social preferences and social norms are in turn explained by indirect evolutionary arguments (e.g., Alger and Weibull 2012; Alger and Weibull 2013). The reason for this separation is that both models are complex enough when the two issues are separated, hence a full formal treatment appears infeasible (see Schelling 1971; Skyrms 2004). The advantage of following the route of the simpler Nash dynamics, as we do in this chapter, is that these models are tractable, and their macro predictions are easy to verify. This approach is also more parsimonious and more general than the preference/norm based approaches (at least as a first step).

This approach to explaining observed behaviors can be successful when the game has more than one equilibrium and/or does not feature a dominant strategy. When there is only one equilibrium (as in L-PGG), one must turn to more complex models to explain why some individuals would ever contribute. But when multiple equilibria exist, and when there is more room for strategizing, then Nash dynamics lead to/from all equilibria to one another, and perturbed Nash dynamics will make predictions about their relative stability.

We shall now provide a simple framework to express several alternative individual-level adjustment dynamics that fall under the category of "Nash dynamics", and we shall use them to try and explain behavioral regularities observed in a number of experiments.

The model. Suppose the same population of players $N = \{1, 2, \ldots, n\}$, by choosing actions from the same finite action set C, repeatedly plays the same symmetric *noncooperative game* in periods $\boldsymbol{T} = \{1, 2, \ldots, t\}$, where each outcome, an action n-tuple $\mathbf{c} = \{c_i\}_{i \in N}$, implies payoff consequences $\boldsymbol{\phi}_{\mathbf{c}} = \{\phi_i(\mathbf{c})\}_{i \in N}$.

Denote by $BR_i(c^t_{-i})$ the *best response* by player i against the actions played by the others in period t, c^t_{-i}. Omitting period superscripts, being a best response means that $\phi_i(BR_i|c_{-i})$, that is, the payoff obtained by i by playing $BR_i(c_{-i})$, is larger than any other payoff obtainable by playing an alternative action (assuming it is unique). Of course, an outcome c is a *Nash equilibrium* if, and only if, $c_i = BR_i(c_{-i})$ for all players.

We assume (along the lines of Kandori, Mailath, and Rob 1993 and Young 1993) that, over time, each player plays $BR_i(c_{-i})$ with probability $(1-\epsilon)$, and all other strategies with some positive probability summing up to ϵ. Note that, as $\epsilon \to 0$, only certain outcomes (called *stochastically stable* states) have positive probability in the long-run distribution of the dynamic (Foster and Young 1990).

Of course, the crucial question for each player is to hypothesize what others' actions will be to decide what one should play oneself. A standard assumption in evolutionary games – and the route taken in Nash (1950) – is to base this hypothesis on information about other players' actions in the past. However, there are several alternative assumptions as to how information from the past is processed, which we will discuss below. There are also a number of differences in sampling from players' pasts, but we shall not address this issue here, assuming that all past actions are perfectly observable (as they are in many experimental settings).

We illustrate these different assumptions using the two-by-two coordination game known as *battle of the sexes*, where each player chooses between two actions, "opera" and "football". Each player receives a payoff of two (zero) from coordination (anti-coordination) on any of the actions, and in addition has an idiosyncratic preference worth an additional payoff of one for one of the two actions (man prefers football, woman prefers opera). Of course, the best response in such a game is always to match the other's action and there are, therefore, two pure strategy Nash equilibria. The important point is that one equilibrium is better for one player, and the other equilibrium is better for the other.

2.1 Basic Nash players

John Nash himself, in his PhD thesis (Nash 1950), formulated the following model of dynamic players (as discussed in Young 2011). We shall call these players *basic Nash players*.

A *basic Nash player* plays

$$BBR^t_i : = BR_i\left(c^{t-1}_{-i}\right),$$

taking as his hypothesis about c^t_{-i} the previous-period observation of others' actions. This model is widely used under various names including myopic best reply. Of course, Nash equilibria are absorbing states of basic Nash play dynamics.

In the context of the battle of the sexes game, if a man observes the woman playing opera this period, he will play opera next period.

2.2 Clever Nash players

Building on the notion of the basic Nash player (as used in Young 1993), Saez-Marti and Weibull (1999) introduce the notion of "cleverness" among agents, leading to so-called *clever Nash players*. A clever Nash player plays a best response against the other player's basic best responses.

More formally, a *clever Nash player* plays

$$CBR_i^t := BR_i\left(BR_{-i}\left(c^{t-1}\right)\right) = BR_i\left(BBR_{-i}^t\right),$$

taking as his hypothesis about c_{-i}^t the current-period basic Nash player best responses, $BBR_j^t = BR_j(c_{-j}^{t-1})$, for all other players j. A clever Nash player predicts, by looking backwards, the current play of basic Nash players, and he can therefore potentially improve his play by unilaterally responding differently to basic best response. Naturally, Nash equilibria also remain absorbing states of clever Nash play dynamics.

In the context of the battle of the sexes game, if a man played football this period he will assume that the woman will play football next period, and therefore play football next period too.

2.3 Forward-looking Nash players

In the spirit of Saez-Marti and Weibull (1999), and adding an element of forward-lookingness, we introduce the notion of "forward-looking cleverness". We shall therefore call such players *forward-looking Nash players*.

A forward-looking Nash player plays FBR_i^t to maximize his next-period payoff. He assumes others will act as basic Nash players next period and will play BBR_j^{t+1} for all j, taking as given FBR_i^t and their BBR_k^t for all $k \neq i, j$. For himself, he intends to play clever Nash CBR_i^{t+1} against BBR_{-i}^{t+1}.

In other words, this means that a forward-looking Nash player predicts the future consequences of his own current-period action on his next-period payoff, and he chooses his own action to maximize his forward-looking payoff. Note that a forward-looking Nash player can therefore consider the consequences of unilateral deviations and further consider the consequences of inducing multilateral deviations. Such a player can move play of the population from one Nash equilibrium to another. In general, however, Nash equilibria need not be absorbing states of forward-looking Nash dynamics. It is worth pointing out that it is an under-explored avenue to link foresightedness with the rich literature on cognitive hierarchy.

In the context of the battle of the sexes game, this means that the man player will play football this period to make the woman player choose football next period, even at the risk of anti-coordination this period.

2.4 Perturbed dynamics

We assume that the order of magnitude of noise added to basic Nash play is higher than that added to clever Nash and forward-looking Nash play. However, noise is added to all types of Nash play. Hence, the process continues to be ergodic (every outcome is reached from every outcome with positive probability), meaning that we can apply ergodic theory and stochastic stability arguments.

Our baseline assumption will be to consider uniform constant deviation rates for each type of agent. An alternative way of introducing noise to underlying Nash play dynamics would be to develop an approach where the probability of an error depends on its cost vis-à-vis the best response. This approach is taken in Blume (1993) and many subsequent contributions, and corresponds to the approach of quantal response equilibrium (McKelvey and Palfrey 1995; McKelvey and Palfrey 1998). Since this chapter is analytical rather than motivated by fitting data, however, we shall go for the simpler (i.e., one parameter) uniform noise assumption in our approach.

3 Social dilemma games

We shall now detail our main game application, called "meritocratic matching games", and introduce the other three classes of games mentioned in the introduction for subsequent comparison.

3.1 Meritocratic matching games

Consider the following *meritocratic matching game* (MMG). All agents in the population $N = \{1, 2, \ldots, n\}$ have to decide simultaneously whether to contribute toward the provision of local public goods (in the sense of a club/team good, Buchanan 1965), choosing an arbitrary amount c_i from some fixed budget B such that $c_i \in [0, B]$. Given the vector of all contributions, c, the population is divided into several groups, of equal size $s < n$, and contributors (freeriders) tend to be matched with contributors (freeriders). We shall call such a matching *meritocratic matching*, and the resulting class of games are the meritocratic matching games.

Meritocratic matching encompasses a range, from no-meritocracy to full-meritocracy, which we shall instantiate as follows. Suppose i.i.d. Gaussian noise, $\epsilon_i \sim (0, \sigma^2)$ with $\sigma \in (0, \infty)$, is added to each actual contribution decision so that, instead of the actual contribution c_i, only the noised contribution, $(c_i + \epsilon_i)$, is observable. Players are then ranked according to $\{(c_i + \epsilon_i)\}_{i \in N}$, from highest to lowest, and groups form composed from this order: the highest s $(c_i + \epsilon_i)$s form group one, etc.

$\beta := 1/\sigma$ represents the *level of meritocracy* in the system.

- For $\beta \to \infty$ (or $\sigma^2 \to 0$), noise vanishes and we approach *full-meritocracy*. In that case, no player contributing less than another can be ranked higher than him, but there is random tie-breaking to ascertain the precise ranking of contributors who contribute the same amount.
- For $\beta \to 0$ (or $\sigma^2 \to \infty$), noise takes over and we approach *no-meritocracy* or random group matching (as, for example, in Andreoni 1988).

The strength of assortativity in our process is expressed by β, an index that can be related to the so-called "index of assortativity" (Bergstrom 2003; Bergstrom 2013; Jensen and Rigos 2014; Nax and Rigos 2016; see also Wright's "F-statistic" 1921; 1922; 1965). It corresponds to the notion of "institutional fidelity" in the sense that some real-world institutions/mechanisms endeavor to be meritocratic but do so imperfectly for a variety of different reasons.

Meritocratic matching in the form of an assortative matching of contributors and freeriders alike creates incentives to contribute to be matched with others doing likewise. Given contribution decisions and groups that form based on these, some *marginal per-capita rate of return*, $r \in (\frac{1}{s}, 1)$, determines the return in each of the n/s local public goods which are shared equally among the agents in each group. Note that there are no payoff transfers between players. Denote by S_i the group in which i is matched. Consequently, i will receive a monetary *payoff* of $\phi_i = B - c_i + r \cdot \sum_{j \in S_i} c_j$.

There is no *a priori* dominant strategy in this game under meritocratic matching (provided β is sufficiently high). Nevertheless, the outcome where all players freeride (contribute zero) is a Nash equilibrium for any value of β. In addition, if σ^2 is not too large and r is sufficiently large, then there exist additional Nash equilibria (Gunnthorsdottir et al. 2010; Nax, Murphy, and Helbing 2014). These are asymmetric outcomes where a large majority, size $m > (n - s)$, of the population contributes fully, while only a marginal minority of players, size $(n - m) < s$, freerides. The exact size of the freeriding minority depends on the game's defining parameters.

Meritocracy has been shown experimentally (Gunnthorsdottir et al. 2010; Rabanal and Rabanal 2014) to effectively implement near-efficient contribution levels at the high Nash equilibrium values. This result has been shown to generalize for the inclusion of noise for general meritocracy levels, theoretically (Nax, Murphy, and Helbing 2014) and experimentally (Nax et al. 2015). It is noteworthy that this mechanism works, in the sense of implementing near-efficient outcomes, with homogenous, narrowly self-regarding (Nash) players.

An issue we have so far left unaddressed is to explain how the population coordinates into play of high equilibria.

3.2 Related games

3.2.1 Linear public goods game

The *voluntary contributions mechanism* (VCM) in the context of a linear public goods game (L-PGG), introduced by Marwell and Ames (1979; see also Isaac, McCue, and Plott 1985 and Isaac and Walker 1988), is used widely to study public goods dilemma games in the behavioral sciences (see reviews by Ledyard 1995 and Chaudhuri 2011). Under the linear public goods game (L-PGG), all agents in the population $S = \{1, 2, \ldots, s\}$ simultaneously decide how much to contribute to a public good, choosing an arbitrary amount c_i from some fixed budget B such that $c_i \in [0, B]$. As before, given all contributions c, together with a *fixed marginal per-capita rate of return* $r \in (\frac{1}{s}, 1)$, the public good is then shared equally among the agents who receive a total payoff of $\phi_i = B - c_i + r \cdot \sum_{j \in S} c_j$. Notice now there is no link between contribution decisions and group matching. Of course, the way to maximize one's payoff, therefore, given any combination of contribution decisions by the others, is to set $c_i = 0$. The unique Nash equilibrium is thus characterized by universal freeriding, meaning non-provision of the public good and lowest collective payoffs.

3.2.2 Step-level public goods

The VCM with decisions restricted to a binary choice of whether to contribute ($c_i = 1$) or to freeride ($c_i = 0$) is a special case of the *step-level public goods game* introduced by Palfrey and Rosenthal (1984), which we shall abbreviate as k-PGG. In k-PGG, agents, again, simultaneously decide whether to contribute or not (now $B = 1$). The public good is then provided and shared equally among the agents if, and only if, at least k agents contribute. If fewer than k agents contribute, payoffs are $\phi_i = 0$ for contributors and $\phi_i = 1$ for free-riders. If at least k agents contribute, payoffs are $\phi_i = s \cdot r - 1$ for contributors and $\phi_i = s \cdot r$ for freeriders. The way to maximize one's own payoff, given exactly $(k - 1)$ contributors among the others, is to contribute. However, for any other number of contributors among the others, a unilateral decision is not pivotal, hence the best response is to freeride. Therefore, the resulting structure of pure-strategy Nash equilibria for $k - PGGs$ with k such that $1 < k < s$ includes equilibria as follows:

(A) there exist multiple asymmetric equilibria, in each of which exactly k players contribute and the others freeride;

(B) there exists a symmetric equilibrium in which all players freeride.

3.2.3 The volunteer's dilemma

An important, and different, special case of $k-PGG$ is when $k = 1$, the *volunteer's dilemma game* (Diekmann 1985), here abbreviated VDG. In VDG, the symmetric equilibrium in which all players freeride falls apart because all players are pivotal when no one volunteers, and the only equilibria in pure strategies are asymmetric such that exactly one player volunteers. A symmetric mixed strategy equilibrium where all players contribute with some positive probability also exists, with the surprising comparative statics that for lower-cost volunteers they will volunteer with a smaller probability in equilibrium than the higher-cost players.

4 Predictions

4.1 Baseline evolutionary predictions

4.1.1 Related games

First, we shall consider the baseline evolutionary predictions for the voluntary contribution mechanism applied to the linear public goods game (L-PGG) and for the step-level public goods game (k-PGG). Universal freeriding is the only stochastically stable equilibrium in L-PGG and k-PGG (Myatt and Wallace 2008). For L-PGG under VCM, the reason is quite simply that there is only one Nash equilibrium, which is the non-cooperative outcome. For k-PGG, the reason is more subtle and is a consequence of the incentive structure. Above the threshold, that is, when there are already sufficiently many cooperators/contributors to provide the public good, freeriding is a better strategy. The same is true two steps below the threshold. Hence, only at the threshold (for a contributor), or one step below (for a defector), is contributing a best reply. A positive chance of miscoordination away from the local attractor of the high Nash equilibrium, say by an ϵ-tremble or a "bad apple" amid the contributors (Myatt and Wallace 2008), therefore, takes players away from the efficient equilibrium (out of the good basin of attraction) toward the freeriding equilibrium (into the bad basin of attraction).

In terms of dynamics over time, evolutionary predictions are quite different between L-PGG and k-PGG, due to the absence of a dominant strategy in k-PGG. In L-PGG, since freeriding is a dominant strategy, we should observe play at or very close to the zero-contributions throughout. By contrast, in k-PGG, we may see initial play at the high-equilibrium outcome, followed by a relatively sharp drop down to freeriding where the process will remain.

Next, we turn to the volunteer's dilemma games (VDG). In VDG, all (pure strategy) Nash equilibria are stochastically stable if the game is symmetric. What happens when there is asymmetry depends on the underlying noise structure? Under uniform deviation rates, all outcomes are stochastically stable. With cost-dependent deviations, the outcome where the player with the lowest provision-cost would volunteer is selected. In terms of evolutionary stability, all Nash equilibria are stable. In the presence of asymmetries amongst players, the evolutionary stable states may additionally include asymmetric mixed-strategy equilibria, and even the welfare-maximizing equilibrium where the lowest-cost player volunteers with probability one may be an evolutionary stable state (He et al. 2014). The dynamics of the evolving game may be quite complex (see Raihani and Bshary 2011; Diekmann and Przepiorka 2015). We may see turn-taking behaviors, lock-in to one specific volunteer, and occasional breakdown or over-volunteering if players follow mixed strategies. A symmetric VDG can have quick lock-in behavior under stochastic fictitious play. If all players have a non-zero probability of volunteering, eventually one of them does, and they are then frozen in that configuration with the unlucky volunteer and everyone else staying out.

4.1.2 Meritocratic matching games

In terms of stability, the class of MMGs is divided roughly into three types:

(A) For games with low meritocracy and/or low rates of return, the freeriding equilibrium is the unique Nash predictions and therefore stable.

(B) For games with intermediate meritocracy and intermediate rate of return, near-efficient equilibria exist, but the freeriding equilibrium is the unique stable equilibrium.

(C) For games with high meritocracy and high rate of return, near-efficient equilibria exist and are stable.

Play of the two extreme cases – (A) and (C) – from the experiment by Nax et al. (2015) is summarized in Figure 1.

In terms of dynamics over time, ex ante, we would expect the following behaviors for the three groups of MMGs. For the first (A), we expect similar dynamics as in L-PGG. Indeed, this is what Figure 1 illustrates. For the second (B), we expect similar dynamics as in k-PGG. Finally, for the third (C), we expect dynamics where initial play could either already be close to a high-equilibrium, or start closer to the freeriding outcome. In the latter case, we would expect a relatively sharp increase in contributions quickly, leading from the freeriding outcome to high-equilibrium cooperation levels. Once at high-equilibrium levels, we would expect the process to remain there with high stability. Figure 1 illustrates experimental play of (C).

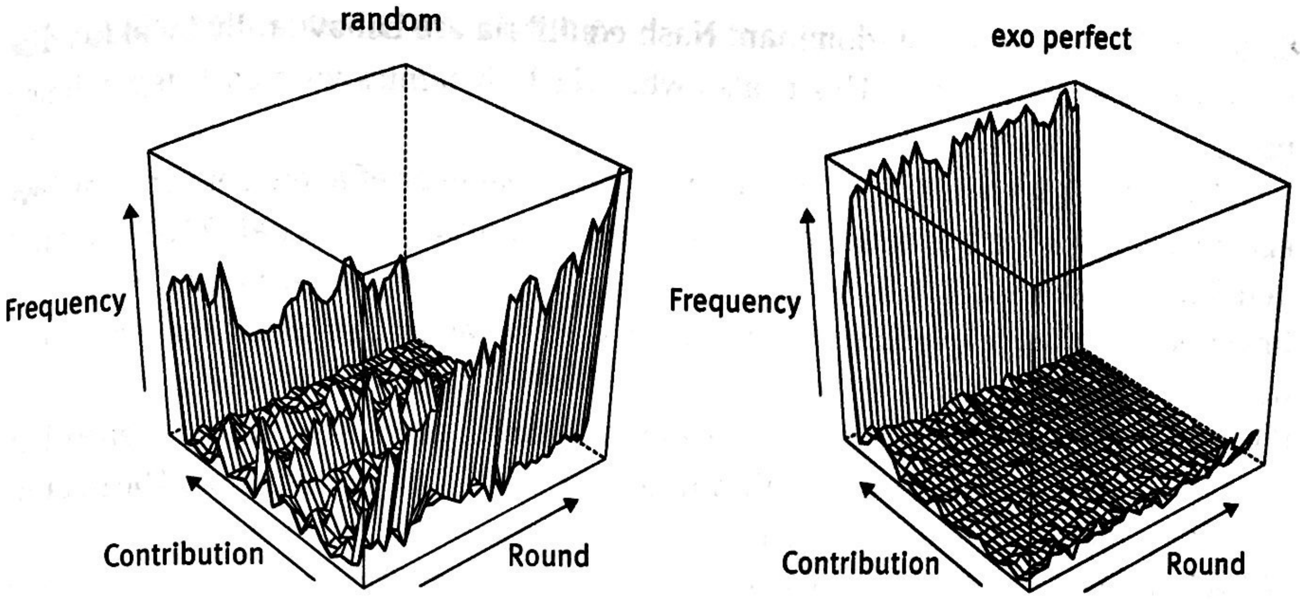

Notes: Treatments varied with respect to the degree of meritocracy in the system, ranging from "no-merit" (random re-matching) to "perfect-merit" (a perfectly meritocratic matching protocol). In the figure, the contribution patterns for the case of no-merit are shown on the left, and for perfect-merit on the right.

Fig. 1: Contribution patterns from a laboratory experiment under meritocratic matching (Source: Nax et al. (2015); kindly produced by S. Balietti).

4.2 Experimental evidence

The experimental evidence for the L-PGG is well known (as reviewed in Ledyard 1995, and more recently in Chaudhuri 2011). Basically, without further mechanisms, contributions start at some intermediate level and decay over time by roughly half the amount every ten rounds under random re-matching (Andreoni 1988), and less when group matching is fixed. The initial contribution pattern is unexpected and is probably best explained by the introduction of additional features such as social preferences (Fischbacher and Gachter 2010; Chaudhuri 2011) or learning (Burton-Chellew, Nax, and West 2015). Nash dynamics cannot explain these high initial levels of cooperation/contribution, but they do explain what happens with iterated interactions.

For k-PGGs, the pattern depends crucially on how many contributors, relative to the population size, are needed. The likelihood that the threshold is met or exceeded is higher for lower thresholds relative to the population size (for important and recent contributions see, for example, Erev and Rapoport 1990; Potters, Sefton, and Vesterlund 2005; Potters, Sefton, and Vesterlund 2007; Gächter et al. 2010a; and Gächter et al. 2010b. A lacuna exists for an up-to-date literature review for k-PPGs).

For VDGs, Diekmann and Przepiorka (2015) discuss much of the relevant behavioral evidence. What is important is that there is evidence of turn-taking. Moreover, the counterintuitive comparative static of the mixed equilibrium where the lowest-cost volunteers volunteer with probabilities lower than the others is behaviorally not confirmed. Instead, they contribute more often, which is more in line with Harsanyi-

Selten logic that the payoff-dominant Nash equilibria are behaviorally focal (as discussed in Diekmann 1993). This is also what Nash dynamics with cost-dependency predict.

For MMGs, there exists evidence that suggests that near-efficient contribution levels are achieved close to what theory predicts (Gunnthorsdottir et al. 2010; Rabanal and Rabanal 2014; Nax, Murphy, and Helbing 2014), and for noisy meritocracy (e.g., imperfect assortative matching). Three aspects of behavioral evidence are especially noteworthy:

(1) The near-efficient equilibrium is the uniquely stable equilibrium, in the sense that outcomes at or close to that equilibrium are played in all experiments. The freeriding equilibrium is never played.
(2) A large fraction of players take turns to freeride in equilibrium. Often, turn-taking functions without, or with very little, loss of equilibrium miscoordination.
(3) Other than in L-PGG, where contributions gradually decline towards the Nash equilibrium over time, there is almost no change in behavior (they may be learning to coordinate and trust each other) in the MMGs. Instead, players play at or very close to the near-efficient equilibrium virtually from the start of the game and continue unabated.

We shall dedicate the remainder of this chapter to explanations of these phenomena.

5 No magic

Kahnemann (1988) speaks of "magic" in the context of market entry games, meaning the observed yet unexplained complex asymmetric equilibria that were successfully coordinated upon in experiments, including turn-taking, "without learning and communication" (Camerer and Fehr 2006:50). This is observed despite individuals acting without clear structure (see Ochs 1999 for a review).

In our MMGs, the asymmetric Nash equilibria are being played too, despite the existence of a simple symmetric Nash equilibrium to which players could resort instead. Hence, in our MMGs, in the sense of Kahnemann (1988:12), it has been suggested that "subjects display more complex coordination and 'magic' than hitherto observed" (Gunnthorsdottir, Vragov, and Shen 2010b). We shall attempt to unravel the magic in this section, based on decision-theoretic foundations that mirror most closely the basic logic of Nash behavior, thus offering a very simple "no-magic argument" for our three phenomena.

5.1 Stability of near-efficiency

First, we shall address the question of why the near-efficient equilibrium is more stable. The reason is simple. All that is needed to jump out of the basin of attraction of the no-contribution equilibrium into that of the near-efficient equilibrium is for two-players to contribute fully; and so as not to exceed Nash predictions, all that is needed are a few players to contribute zero.

The latter can be explained by basic Nash play and by clever Nash play, and the former by forward-looking Nash play. The reason is that the basic best response to all other players contributing fully is to contribute zero, while the forward-looking best response is to contribute fully once a few players contribute fully, which will subsequently lead to near-efficient contributions. Mistakes by basic Nash players, therefore, together with clever and forward-looking Nash play by just a few agents for each category, explains why the asymmetric, near-efficient equilibrium will be played quickly. Without forward-looking Nash players, the efficient outcome could also emerge after some time with some noisy players and best response.

5.2 Turn-taking

Next, we shall address the phenomenon of turn-taking. Suppose we start the process off in the near-efficient equilibrium, assuming it exists. If so, then predicting the tremble of a contributor will sway a freerider playing clever best response to contribute fully. Similarly, out of equilibrium: if too few freeriders exist vis-à-vis the near-efficient equilibrium, then a clever Nash player currently freeriding will switch to contributions, expecting the basic Nash players to implement the freeriding strategy. Similarly, if too many freeriders exist vis-à-vis the near-efficient equilibrium, then a clever Nash player currently contributing will switch to freeriding, expecting the basic Nash players to implement the contribution strategy.

Low-probability mistakes by basic Nash players, therefore, together with clever Nash play by one or two agents, could explain turn-taking.

5.3 No learning

The reason why no learning is required to account for the near-efficient equilibrium virtually from the beginning is because most players can reactively play a basic Nash best response. In fact, this basic Nash play is what allows the clever and forward-looking players to coordinate between different equilibria. In addition, for the near-efficient Nash equilibrium, there are no mixed motives between individual and group incentives for the large majority of players in equilibrium who will contribute fully. It is best for them to play it and it is best for the collective that it is played.

5.4 Predictions for related games

In L-PGG under VCM, our model predicts zero contributions throughout the iterated games. Evidence of intermediate contributions would imply astonishingly high deviation rates, which would ebb as contributions decline. This explanation is unsatisfactory, and alternative explanations are required.

In k-PGGs, our model predicts that (provided sufficiently many clever and forward-looking players exist) the threshold would be reached or even exceeded. In fact, it offers a simple explanation of the phenomena, suggesting that early contributors are clever or forward-looking, which is very much in line with the explanations that have been proposed.

In VDGs, our predictions are especially interesting. For example, a forward-looking Nash player will never volunteer unless he expects no other player to do so, or if currently there is over-volunteering and he expects a backlash. Similarly, a clever Nash player would volunteer if currently there is over-volunteering, but never if there is no volunteer already. Basic Nash players, on the other hand, react to the current market pressure. In an asymmetric volunteer's dilemma, this implies that the strong players will manage to avoid volunteering if they are clever or forward-looking, but not if they are basic Nash players.

6 Concluding remarks

Many interactions are such that Nash equilibria predictions are highly sensitive to the exact assumptions we make about the agents' utilities and about their beliefs about the other players. The standard linear public goods game is one such example. Based on pure material self-interest, we could not explain why subjects in laboratory experiments consistently contribute positive amounts. Hence, one needs to turn to more complex models involving bounded rationality, social preferences, reciprocity, etc. The exact modeling assumptions will then be crucial for predicting how much is contributed under the VCM and at what time.

Other interactions are different. In some games, no matter what assumptions we make about the agents' utility functions and about their beliefs, one and the same type of outcome is generally predicted. The meritocratic matching game was such an interaction. In situations like this, models of perturbed Nash dynamics make robust and accurate predictions, while the many explanations from other related games (such as the standard linear public goods games) do not.

The aim of this chapter was to illustrate how predictions about (non-) cooperative behavior can be made on the basis of various perturbed Nash dynamics in the context of social dilemma games, particularly those involving meritocratic matching. One way to reconcile our findings that no single explanation could account for the evidence across social dilemmas is to conclude that different game structures and institutions influence different preferences and foster the emergence of different beliefs.

Our main conclusion is that simple noise-driven learning models can explain a great deal of what we observe in various social dilemma games, although this does not apply to all the stylized facts that are commonly observed in experimental play of linear public goods games. The principle of parsimony would dictate employing these simpler learning based approaches before resorting to more involved theoretical approaches based on preferences and beliefs at the individual level. At first blush, this conclusion may appear to encroach on the explanatory power of preferences and beliefs in understanding decision-making in games. On the contrary, by hierarchically accounting for phenomena, and first explaining as much as possible with noise and imperfect players, what remains may then be better explained by more complex and idiosyncratic preferences and beliefs models.

Bibliography

[1] Alger, Ingela, and Jörgen W. Weibull. 2012. "A generalization of Hamilton's rule – Love others how much?" *Journal of Theoretical Biology* 299:42–54.

[2] Alger, Ingela, and Jörgen W. Weibull. 2013. "Homo Moralis – Preference Evolution Under Incomplete Information and Assortative Matching." *Econometrica* 81(6):2269–2302.

[3] Andreoni, James. 1990. "Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving." *The Economic Journal* 100(401):464–477.

[4] Axelrod, Robert. 1984. *The Evolution of Cooperation*. New York: Basic Books.

[5] Axelrod, Robert, and William D. Hamilton. 1981. "The Evolution of Cooperation." *Science* 211(4489):1390–1396.

[6] Bergin, James, and Bart L. Lipman. 1996. "Evolution with State-Dependent Mutations." *Econometrica* 64(4):943–956.

[7] Bergstrom, Theodore C. 2003. "The Algebra of Assortative Encounters and the Evolution of Cooperation." *International Game Theory Review* 5(3):211–228.

[8] Bergstrom, Theodore C. 2013. "Measures of Assortativity." *Biological Theory* 8(2):133–141.

[9] Blume, Lawrence. E. 1993. "The Statistical Mechanics of Strategic Interactions." *Games and Economic Behavior* 5(3):387–424.

[10] Blume, Lawrence. E. 2003. "How noise matters." *Games and Economic Behavior* 44(2):251–271.

[11] Buchanan, James M. 1965. "An Economic Theory of Clubs." *Economica* 32(125):1–14.

[12] Burton-Chellew, Maxwell N., Heinrich H. Nax, and Stuart A. West. 2015. "Payoff-based learning explains the decline in cooperation in public goods games." *Proceedings of the Royal Society B*. doi:10.1098/rspb.2014.2678.

[13] Bush, Robert R., and Frederick Mosteller. 1955. *Stochastic Models for Learning*. New York: Wiley & Sons.

[14] Camerer, Colin F., and Ernst Fehr. 2006. "When Does "Economic Man" Dominate Social Behavior?" *Science* 311(5757):47–52.

[15] Chaudhuri, Ananish. 2011. "Sustaining cooperation in laboratory public goods experiments: a selective survey of the literature." *Experimental Economics* 14(1):47–83.

[16] Darwin, Charles. 1871. *The Descent of Man, and Selection in Relation to Sex*. London: John Murray.

[17] Dawes, Robyn M., and Richard Thaler. 1988. "Anomalies: Cooperation." *Journal of Economic Perspectives* 2(3):187–197.

[18] Dawkins, Richard. 1976. *The Selfish Gene*. Oxford: Oxford University Press.

[19] Diekmann, Andreas. 1985. "Volunteer's Dilemma." *Journal of Conflict Resolution* 29(4):605–610.

[20] Diekmann, Andreas. 1993. "Cooperation in an asymmetric Volunteer's dilemma game theory and experimental evidence." *International Journal of Game Theory* 22(1):75–85.

[21] Diekmann, Andreas, and Wojtek Przepiorka. 2015. "Punitive preferences, monetary incentives and tacit coordination in the punishment of defectors promote cooperation in humans." *Scientific Reports* 5:17–52.

[22] Doebeli, Michael, and Christoph Hauert. 2005. "Models of cooperation based on the Prisoner's Dilemma and the Snowdrift game." *Ecology Letters* 8(7):748–766.

[23] Erev, Ido, and Amnon Rapoport. 1990. "Provision of Step-Level Public Goods: The Sequential Contribution Mechanism." *Journal of Conflict Resolution* 34(3):401–425.

[24] Erev, Ido, and Alvin E. Roth. 1998. "Prediction How People Play Games with Unique, Mixed Strategy Equilibria." *American Economic Review* 88(4):848–881.

[25] Eshel, Ilan, Larry Samuelson, and Avner Shaked. 1998. "Altruists, Egoists, and Hooligans in a Local Interaction Model." *American Economic Review* 88(1):157–179.

[26] Fischbacher, Urs, and Simon Gächter. 2010. "Social Preferences, Beliefs, and the Dynamics of Free Riding in Public Good Experiments." *American Economic Review* 100(1):541–556.

[27] Fletcher, Jeffrey A., and Michael Doebeli. 2009. "A simple and general explanation for the evolution of altruism." *Proceedings of the Royal Society B* 276(1654):13–19.

[28] Fletcher, Jeffrey A., and Michael Doebeli. 2010. "Assortment is a more fundamental explanation for the evolution of altruism than inclusive fitness or multilevel selection: reply to Bijma and Aanen." *Proceedings of the Royal Society B* 277(1682):677–678.

[29] Foster, Dean, and Peyton Young. 1990. "Stochastic Evolutionary Game Dynamics." *Theoretical Population Biology* 38(2):219–232.

[30] Gächter, Simon, Daniele Nosenzo, Elke Renner, and Martin Sefton. 2010a. "Sequential vs. simultaneous contributions to public goods: Experimental evidence." *Journal of Public Economics* 94(7–8):515–522.

[31] Gächter, Simon, Daniele Nosenzo, Elke Renner, and Martin Sefton. 2010b. "Who makes a good leader? Cooperativeness, optimism, and leading-by-example." *Economic Inquiry* 50(4):953–967.

[32] Gunnthorsdottir, Anna, Roumen Vragov, Stefan Seifert, and Kevin McCabe. 2010. "Near-efficient equilibria in contribution-based competitive grouping." *Journal of Public Economics* 94(11–12):987–994.

[33] Gunnthorsdottir, Anna, Roumen Vragov, and Jianfei Shen. 2010b. "Tacit coordination in contribution-based grouping with two endowment levels." Pp. 13–75 in *Research in Experimental Economics*. Vol. 13, *Charity with Choice*, edited by R. M. Isaac, and D. Norton. Bingley: Emerald Group Publishing Limited.

[34] Hamilton, William D. 1964a. "Genetical evolution of social behavior I." *Journal of Theoretical Biology* 7(1):1–16.

[35] Hamilton, William D. 1964b. "Genetical evolution of social behavior II." *Journal of Theoretical Biology* 7(1):17–52.

[36] Hardin, Garrett. 1968. "The Tragedy of the Commons." *Science* 162(3859):1243–1248.

[37] Harsanyi, John C., and Reinhard Selten. 1988. *A General Theory of Equilibrium Selection in Games*. Cambridge, MA: The MIT Press.

[38] He, Jun-Zhou, Rui-Wu Wang, and Yao-Tang Li. 2014. "Evolutionary Stability in the Asymmetric Volunteer's Dilemma." PLoS ONE 9(8):e103931.

[39] Helbing, Dirk. 1992. "Interrelations between Stochastic Equations for Systems with Pair Interactions." *Physica A: Statistical Mechanics and its Applications* 181(1–2):29–52.

[40] Helbing, Dirk. 1996. "A Stochastic Behavioral Model and a 'Microscopic' Foundation of Evolutionary Game Theory." *Theory and Decision* 40(2):149–179.

[41] Helbing, Dirk. 2010. "The future of social experimenting." *PNAS* 107(12):5265–5266.

[42] Helbing, Dirk. 2015. *The Automation of Society Is Next: How to Survive the Digital Revolution*. North Charleston, SC: CreateSpace Independent Publishing Platform.

[43] Isaac, Mark R., and James M. Walker. 1988. "Group Size Effects in Public Goods Provision: The Voluntary Contributions Mechanism." *Quarterly Journal of Economics* 103(1):179–199.

[44] Isaac, Mark R., Kenneth F. McCue, and Charles R. Plott. 1985. "Public goods provision in an experimental environment." *Journal of Public Economics* 26(1):51–74.

[45] Jensen, Martin K., and Alexandros Rigos. 2014. *Evolutionary Games with Group Selection*. Working Paper No. 14/9, University of Leicester.

[46] Kahneman, Daniel. 1988. "Experimental Economics: A Psychological Perspective." Pp. 11–20 in *Bounded Rational Behavior in Experimental Games and Markets*, edited by R. Tietz, W. Albers, and R. Selten. Berlin: Springer.

[47] Kandori, Michihiro, George J. Mailath, and Rafael Rob. 1993. "Learning, Mutation, and Long Run Equilibria in Games." *Econometrica* 61(1):29–56.

[48] Van de Kragt, Alphons J. C., John M. Orbell, and Robyn M. Dawes. 1983. "The Minimal Contributing Set as a Solution to Public Goods Problems." *American Political Science Review* 77(1):112–122.

[49] Ledyard, John O. 1995. "Public Goods: A Survey of Experimental Research." Pp. 111–194 in *Handbook of Experimental Economics*, edited by J. H. Kagel, and A. E. Roth. Princeton, NJ: Princeton University Press.

[50] Mäs, Michael, and Dirk Helbing. 2014. "Noise in behavioral models can improve macro-predictions when micro-theories fail." Preprint.

[51] Mäs, Michael, and Heinrich H. Nax. 2015. "A Behavioral Study of "Noise" in Coordination Games." *Journal of Economic Theory* 162:195–208.

[52] Marwell, Gerald, and Ruth E. Ames. 1979. "Experiments on the Provision of Public Goods. I. Resources, Interest, Group Size, and the Free-Rider Problem." *American Journal of Sociology* 84(6):1335–1360.

[53] Maynard Smith, John. 1987. "The Theory of Games and the Evolution of Animal Conflicts." *Journal of Theoretical Biology* 47(1):209–221.

[54] Maynard Smith, John, and George R. Price. 1973. "The logic of animal conflict." *Nature* 246(5427):15–18.

[55] McKelvey Richard D., and Thomas R. Palfrey. 1995. "Quantal response equilibria for normal form games." *Games and Economic Behavior* 10(1):6–38.

[56] McKelvey, Richard D., and Thomas R. Palfrey. 1998. "Quantal response equilibria for extensive form games." *Experimental Economics* 1(1):9–41.

[57] Myatt, David P., and Chris Wallace. 2008. "An evolutionary analysis of the volunteer's dilemma." *Games and Economic Behavior* 62(1):67–76.

[58] Nash, John. 1950. "Non-cooperative games." Ph.D. thesis, Princeton University.

[59] Nash, John. 1951. "Non-cooperative games." *The Annals of Mathematics* 54(2):286–295.

[60] Nax, Heinrich H., Stefano Balietti, Ryan O. Murphy, and Dirk Helbing. 2015. *Meritocratic Matching Can Dissolve the Efficiency-Equality Tradeoff: The Case of Voluntary Contributions Games*. Mimeo (SSRN 2604140).

[61] Nax, Heinrich H., Ryan O. Murphy, and Dirk Helbing. 2014. "Stability and welfare of 'merit-based' group-matching mechanisms in voluntary contribution game." Risk Center working paper, ETH Zurich.

[62] Nax, Heinrich H., and Matjaž Perc. 2015. "Directional Learning and the Provisioning of Public Goods." *Scientific Reports* 5(8010). doi:10.1038/srep08010.

[63] Nax, Heinrich H., Matjaž Perc, Attila Szolnoki, and Dirk Helbing. 2015a. "Stability of cooperation under image scoring in group interactions." *Scientific Reports* 5:12145.

[64] Nax, Heinrich H., and Alexandros Rigos. 2016. "Assortativity Evolving from Social Dilemmas." *Journal of Theoretical Biology* 395:194–203.

[65] Nowak, Martin A., and Robert M. May. 1992. "Evolutionary Games and Spatial Chaos." *Nature* 359(6398):826–829.

[66] Ochs, Jack. 1999. "Coordination in market entry games." Pp. 143–172 in *Games and Human Behavior: Essays in Honor of Amnon Rapoport*, edited by D. Budescu, I. Erev, and R. Zwick. New York: Erlbaum.

[67] Olson, Mancur. 1965. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Cambridge, MA: Harvard University Press.

[68] Ostrom, Elinor. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge: Cambridge University Press.

[69] Ostrom, Elinor, James Walker, and Roy Gardner. 1994. *Rules, Games, and Common-Pool Resources*. Michigan: University of Michigan Press.

[70] Ostrom, Elinor. 2000. "Collective Action and the Evolution of Social Norms." *Journal of Economic Perspectives* 14(3):137–158.

[71] Ostrom, Elinor, and James Walker. 2003. *Trust and Reciprocity: Interdisciplinary Lessons from Experimental Research*. New York: Russell Sage Foundation.

[72] Ostrom, Elinor. 2005. *Understanding Institutional Diversity*. Princeton, NJ: Princeton University Press.

[73] Palfrey, Thomas R., and Howard Rosenthal. 1984. "Participation and the provision of discrete public goods: a strategic analysis." *Journal of Public Economics* 24(2):171–193.

[74] Potters, Jan, Martin Sefton, and Lise Vesterlund. 2005. "After You – Endogenous Sequencing in Voluntary Contribution Games." *Journal of Public Economics* 89(8):1399–1419.

[75] Potters, Jan, Martin Sefton, and Lise Vesterlund. 2007. "Leading-by-example and signaling in voluntary contribution games: an experimental study." *Economic Theory* 33(1):169–182.

[76] Rabanal, Jean P., and Olga A. Rabanal. 2014. *Efficient Investment via Assortative Matching*: A laboratory experiment. Mimeo.

[77] Raihani, Nichola J., and Redouan Bshary. 2011. "The evolution of punishment in n-player public goods games: a volunteer's dilemma." *Evolution* 65(10):2725–2728.

[78] Rapoport, Amnon, and Ramzi Suleiman. 1993. "Incremental Contribution in Step-Level Public Goods Games with Asymmetric Players." *Organizational Behavior and Human Decision Processes* 55(2):171–194.

[79] Roth, Alvin E., and Ido Erev. 1995. "Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term." *Games and Economic Behavior* 8(1):164–212.

[80] Saez-Marti, Maria, and Jörgen W. Weibull. 1999. "Clever Agents in Young's Evolutionary Bargaining Model." *Journal of Economic Theory* 86(2):268–279.

[81] Schelling, Thomas. 1971. "On the ecology of micromotives." *Public Interest* 25:61–98.

[82] Skyrms, Brian. 2004. *The Stag Hunt Game and the Evolution of Social Structure*. Cambridge: Cambridge University Press.

[83] Taylor, Peter D., and Leo B. Jonker. 1978. "Evolutionary stable strategies and game dynamics." *Mathematical Biosciences* 40(1–2):145–156.

[84] Thorndike, Edward L. 1898. *Animal Intelligence: An Experimental Study of the Associative Processes in Animals*. New York: Macmillan.

[85] von Neumann, John, and Oskar Morgenstern. 1944. *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.

[86] Weibull, Jorgen W. 1995. *Evolutionary Game Theory*. Cambridge, MA: MIT Press.

[87] West, Stuart A., Ashleigh S. Griffin, and Andy Gardner. 2007. "Evolutionary Explanations for Cooperation." *Current Biology* 17(16):R661–R672.

[88] West, Stuart A., Claire El Mouden, and Andy Gardner. 2011. "Sixteen common misconceptions about the evolution of cooperation in humans." *Evolution and Human Behavior* 32(4):231–262.

[89] Wright, Sewall. 1921. "Systems of Mating." *Genetics* 6(3):111–178.

[90] Wright, Sewall. 1922. "Coefficients of inbreeding and relationship." *American Naturalist* 56(645):330–338.

[91] Wright, Sewall. 1965. "The interpretation of population structure by F-statistics with special regard to systems of mating." *Evolution* 19(3):395–420.

[92] Xie, Yu, Siwei Cheng, and Xiang Zhou. 2015. "Assortative mating without assortative preference." *PNAS* 112(19):5974–5978.

[93] Young, H. Peyton. 1993. "The evolution of conventions." *Econometrica* 61(1):57–84.

[94] Young, H. Peyton. 2008. "Social norms." In *The New Palgrave Dictionary of Economics*, 2nd edn., edited by S. N. Durlauf, and L. E. Blume. London: Palgrave Macmillan

[95] Young, H. Peyton. 2009. "Learning by trial and error." *Games and Economic Behavior* 65(2):626–643.

[96] Young, H. Peyton. 2011. "Commentary: John Nash and evolutionary game theory." *Games and Economic Behavior* 71(1):12–13.

[97] Young, H. Peyton. 2015. "The Evolution of Social Norms." *Annual Review of Economics* 7:359–387.