

Choosing to Cooperate: Modelling Public Goods Games with Team Reasoning

15th January 2020

Abstract

This paper presents an agent-based model of *team reasoning* in a social dilemma game. Starting from the conundrum of empirically high levels of cooperation in dilemma games, contradicting traditional utility maximisation assumptions of game theory, Bacharach (1999, 2006) developed a theory of *team reasoning*. The idea behind team reasoning is that agents do not try to maximise their own utility but make choices as part of a team. This paper presents a model of preference convergence, mirroring adaptation dynamics of team reasoning. It describes an agent-based model simulating a repeated public goods game between a designated set of agents, a team. In the model agents have a probability to choose cooperation or defection, adjusting this preferences in the face of the revealed preferences of other players. The model is a classic binary choice model mapping an individual's preference for cooperation onto the binary behaviour choice of cooperation and defection. Preferences are updated in reaction to the behaviour choices of the team. Starting from simple stated preferences, the model implements a reframing of utility maximisation as applying to a group rather than an individual, modelling the importance of social interaction for individual preferences and the dependency of choice on social context. Results show that team reasoning, as implemented here, can explain high levels of cooperation found in the real world resulting from a wide range of settings. It also shows that team reasoning, as implemented here, is not a 'sucker' strategy except when adaptation rates are very slow. This paper demonstrates how agent-based models can be used to examine the role of social contexts for individual decision making.

1 Introduction

In this paper we will look in more detail at how agent-based modelling (ABM) can contribute to two aspects of choice modelling, the question of modelling how individuals actually make choices and of exploring macro outcomes of aggregate individual choices under social influence. The ABM presented in this paper models the importance of an agent's social context on the choice to cooperate or

defect in a public goods game (PGG). The model implements team reasoning akin to Bacharach’s theory of *circumspect team reasoning* (Bacharach, 1999, 2006). Experimental studies strongly support the existence of team reasoning strategies (Colman et al., 2008a,b).

The model departs from the individualist profit maximisation perspective of traditional game theory that sees cooperation as resulting from social order mechanisms such as social norms, reputation or sanctioning. Instead of maximising ones individual payoffs, (some) agents interpret a game or situation as collective, believing others do the same (Elsenbroich and Verhagen, 2016). The model presented here is a model of a PGG but rather than modelling an individual’s focus on utility or payoff the focus is on behaving as part of a team. It explores the influence of preference convergence on levels of individual wealth and social inequality.

When it comes to modelling how individuals actually make choices, the number of theories is almost endless. Different theories can be mapped on a spectrum of universality. Some theories are very specific, taking into account personality differences or an individual’s life course (e.g. Kassarian, 1971; Ham et al., 2009), focus on detailed contexts, looking at socio economic variables such as gender, class and ethnicity (e.g. Dekkers et al., 2000) or the influence of social networks (e.g. Easley et al., 2010). On the other end of the spectrum are universal theories that capture choices with a single principle, e.g. utility maximisation (e.g. Hooff et al., 2001).

The increase in universality goes hand in hand with a loss of heterogeneity of individuals. For formal modelling at scale this universality used to be essential as increasing heterogeneity automatically results in models too complex to solve analytically. The loss of heterogeneity, however, means that potentially salient information is lost. This is a particular problem for modelling the macro or societal consequences of individuals’ choices. When it comes to predicting the future state of a system such as a society, small differences in the individual components, the agents, can lead to big differences in overall system behaviour (Ormerod, 2011).

The question of the consequences at the social, societal or macro level resulting from the aggregation of individual choices is the second dimension we are looking at in this paper. At the more abstract level these are questions about the sustainability, equality and equity of societies (e.g. Luebker, 2014). At a more concrete level there might be questions about how choices, such as education, influence the life course (e.g. Cardak, 2004) or what might lead to migration decisions (Klabunde and Willekens, 2016). Understanding the societal dimensions of individual choices is at the heart of many social sciences, nowhere more so than in economics. Economics built on the choice model of rationality has had to recognise that irrationality and bias are integral parts of human decision making (Kahneman, 2011). This heterogeneity detracts from the simplicity of the models but is paramount for understanding how individual choices aggregate.

Agent-based modelling (see section 3) is a method that allows for formal and systematic modelling of interactions of heterogeneous individuals. Whilst ABM

cannot find out how people actually make choices (this is a question reserved for empirical research) ABM is a method that can implement heterogeneity, allowing for modelling of less abstract, more context dependent and concrete ways of making choices than traditional mathematical approaches. Through this ABM has the ability to model macro social outcomes of social influences and interactions. Klabunde and Willekens argue strongly for the use of ABM, as “Agent-based modelling is the only method that allows for the explicit modelling of social interaction and the social networks that result from it.” (Klabunde and Willekens, 2016, p.75). By modelling many agents over time, interacting with each other and the environment in which they are situated, ABM can simulate macro phenomena emergent from the micro interactions. Furthermore, through modelling individual agents over time interacting with the macro level, ABM can also model how the macro level influences the individuals (*immersion*, cf. Conte et al. (2014)).

The ABM discussed in this paper models individuals’ choices to cooperate or defect in a social dilemma setting, depending on the group context of the agent. It implements a particular theory of choice, team reasoning (Bacharach, 1999, 2006), as an explanation for high levels of cooperation and models the macro consequences of team reasoning in different social contexts. In the model the idea of team reasoning is operationalised by an agent observing the cooperation levels within their group or team and adjusting their personal likelihood of cooperation in the next iteration.¹

The purpose of this model is theoretical exploration of team reasoning as an alternative theory of choice to traditional individual utility maximisation. The model contributes to the literature on cooperation in collective dilemmas by looking at cooperation dynamics where individuals make the decision whether to cooperate or not in relation to their membership of a collective or team. Rather than focussing on utility calculations it focusses on convergence dynamics of cooperation probabilities in a team. The questions the model is trying to tackle are outcomes of individual wealth and population inequality for different initial group compositions, as is done in traditional social dilemma models. The dynamics driving the model are, however, interaction convergence dynamics similar to traditional opinion dynamics models.

The paper is structured as follows. In Section 2 we discuss some solutions to the problem of cooperation in social dilemmas, in particular team reasoning. In Section 3 we introduce agent-based modelling as a method for modelling the choice to cooperate or defect. In Section 4 we describe the model, presenting the analysis and results in Section 5. Section 6 draws conclusions and discusses future work.

¹This is akin to the notion of team reasoning developed in (cf. Bacharach, 1999, 2006) which sees team reasoning as an unconscious reaction to a social situation, in contrast to Sugden (2011) explicitly negotiated version.

2 Explaining Cooperation

The problem for cooperation in social dilemma settings is the temptation for each individual of free-riding. In situations where collective effort is required to allocate limited resources, a player can do better for themselves by using the resource without contributing. In case of too many free-riders the collective effort collapses and everybody loses out. Everyday examples are the colleague who does not attend meetings or the flatmate who does not contribute to housework. More formal examples are the tragedy of the commons (ToC) (Hardin, 1968), in which a public resource is used at different rates by multiple agents, or a public goods game (PGG), in which contributions to a public good are shared out (e.g., Brañas-Garza and Espinosa, 2011).

The analysis of cooperation is interested in the macro outcomes, whether and under what circumstances cooperation establishes or ceases. When it comes to how choices are made, the modelling of (repeated) cooperation dilemmas has mostly been analysed from the vantage point of agents that make their choice to cooperate or defect according to fairly simple strategies. Within the discipline of game theory this is understandable as it is itself part of economics which generally relies on relatively simple decision models. However, there has been a tendency to analyse quite concrete social phenomena using principles of utility maximisation of game theory. For example games are being used to model social phenomena such as the choice of spouse, migration, decisions about education and job markets or even Mafia extortion (Becker, 2017; Gambetta, 2000).

The main problem in the game theoretic analysis of cooperation games is that empirical levels of cooperation are much higher than the theory predicts. Cooperation is surprising in a framework of theories of individual (evolutionary) advantage where payoffs for defectors are higher than for cooperators. This is particularly clear in stylised collective dilemmas such as the ToC or a PGG. In the real world, however, many collective situations work, e.g., flat shares, coops, public broadcasting and tax systems. Also, in experiments of social dilemma situations, whether the prisoner’s dilemma or collective games such as a PGG, defection rates are around 30%: unexpectedly low given an individual utility maximisation framework (Fischbacher et al., 2001; Poteete et al., 2010).

That levels of cooperation are surprisingly high can be seen as a clash between the highly universal theory of game theoretic choice and experimental findings with real human beings. There have been many different approaches to the problem of collective dilemmas (see, e.g., Janssen and Ostrom, 2006, as an overview). Experimental studies investigate the influence of reciprocity, fairness and greed in public goods dilemmas (Fehr and Gächter, 2000; Eek and Biel, 2003) as well as structural features of communication and team heterogeneity (Skatova and Ferguson, 2013). Burton-Chellew and West (2013) found that prosocial preferences are not predictors for behaviour in PGG. Overall, experimental game theory finds high levels of cooperation, between 50 and 70% (Burton-Chellew et al., 2016). Conditional cooperators, i.e. those that cooperate if they experience cooperation of others, consistently make up around 50% of the population, free-riders 25%, and the remaining 25% have a mix of strategies including some

unconditional cooperators (Burton-Chellew et al., 2016).

Within cooperation research a host of theoretical starting points have been advanced to bridge this clash. Many different ideas have been advanced to explain this discrepancy between the theoretical prediction of high and the empirically found low levels of defection. Some explanations invoke social mechanisms such as sanctioning or ostracism (Balliet et al., 2011), some show the evolutionary advantage of reciprocity (direct or indirect) in the long run (Boyd and Richerson, 1989; Nowak and Sigmund, 1998; Poteete et al., 2010) and some look at cognitive mechanisms such as Cognitive Hierarchy Theory, (e.g., Camerer et al., 2004). Bicchieri (2006) investigates the viability of social norms to explain cooperation in strategic interactions. Ostrom’s work tackles explicitly collective dilemmas through the perspective of social norms for enforcing cooperation (Ostrom, 1990; Poteete et al., 2010; Ostrom, 2014). Gintis and Bowles (2011) argue for the evolutionary strength of cooperation using reciprocity, cultural evolution and group selection. All of the above approaches stipulate that the choice to defect is constrained by reducing the utility of defection.

This is different in the theory that inspired the model in this paper, the theory of *team reasoning* in (Bacharach, 1999, 2006). Team reasoning is developed as an explanation of the high levels of cooperation found in social dilemma experiments without giving up utility maximisation. In team reasoning, rather than amending the individual payoff function, the unit to which utility maximisation is applied changes from the individual to the team. Some situations are recognised by players as collective (or team) situations, demanding cooperation *whatever the consequences for the individual*. The prisoner’s dilemma is such a situation, where the change of strategy by both players to cooperate as a team brings about collectively superior outcomes; collective dilemmas are another example. Team reasoning doubts the superiority of the Nash equilibrium. Whilst for each individual player the *temptation payoff* is the profit maximising outcome, the joint best outcome for all players is the cooperation payoff. In team reasoning, the assumption of rationality of both (all) players becomes “if both (all) of us are rational, *we* will cooperate.” The important difference when comparing team reasoning to other choice mechanisms is that rather than changing the utility function, e.g. by the introduction of sanctions, altruism or fairness norms, team reasoning changes who the utility function is applied to, the group or team, rather than the individual. Rationality is still inherently about utility maximisation but not individual utility maximisation. As Colman and Gold (2018, p. 1774) point out:

“The change of agency is both subtle and radical; it involves a concept of group agency in which rationality is transferred from individual players’ actions to the joint action of the group of players as a whole, although the decisions are ultimately taken by individuals.”

That means team reasoning keeps the universality of utility maximisation but contextualises it through the social situation an individual finds itself in. This has the consequence that agents can be suckers, i.e. cooperate whilst other

players defect. Bacharach calls this *circumspect* team reasoning, where agents are weary and, once they realise the others in the team do not interpret the situation as collective, revert to defection.²

Experimental studies have found evidence for genuine team reasoning in comparison of inter-social mechanisms (Colman et al., 2008a,b), motivating us to extend the computational research into cooperation in collective dilemmas to also implement collective kinds of reasoning. What is important in team reasoning is not utility calculations but a dynamic of adaptation to a situation misinterpreted as a team, here modelled through the application of opinion dynamics to the cooperation probabilities.

3 Agent-based Modelling

ABM has been successfully used in disparate fields such as geography, criminology, economics, organisation research, epidemiology and sociology (cf., Squazzoni, 2010; Heppenstall et al., 2011; Liu and Eck, 2008; Ostrom, 2014; Watts and Gilbert, 2013; Badham et al., 2018; Troitzsch, 2018). An agent-based model (ABM) is a computer program that models individual agents (people, households, organisations, etc.) and their interactions with each other and a given environment (resources, information, networks etc). Through modelling interactions over time, ABM can generate patterns representing emergent social phenomena, modelling the development of a (social) system over time (Gilbert, 2008). Some particular advantages of agent-based modelling are:

Modelling Processes: ABM is a simulation method replicating processes of a target system. This strongly differentiates ABM from many other social science methods which are often static. The possibility to model processes has two consequences for modelling choice. It can model the temporal aspects of individual choices and the emergence of social phenomena over time and. Agent-based modelling can account for a prolonged, multi-step decision process as for example advocated in Janis and Mann (1977), which stresses the importance of accounting for the discrete steps of decision making.

Incomplete Knowledge: Agents in an ABM make decisions based on rules and heuristics. This means that ABM can easily model incomplete knowledge, in contrast with analytic methods. Whilst rational choice and game theory are important foundations of ABM (Axelrod and Hamilton, 1981), different ways of making choices can be implemented in ABM, using empirically supported decision making rules extracted from qualitative research (e.g. Yang and Gilbert,

²This is very different from accounts of team reasoning developed by Sugden. According to Sugden (2003, 2011), agents being suckers must be explicitly ruled out: people only team reason if they are “mutually assured” of payoff, i.e., assured that everyone else in the team will cooperate. Mutual agreement to cooperate is a prerequisite to the application of team reasoning in Sugden’s theory.

2008).

Feedback and Adaptation: By agents interacting with each other and their environment, agents change their environment, leading to a feedback between agent decision making and emerged macro phenomena (Conte et al., 2014). ABM can thus model behaviour change as a result and driver of system change.

Decision Contexts: One of the main features of ABM is the ability to model decision making in context. Agents in ABM take their situation into account when making decisions. The situation can comprise an agent's current environment as well as personal or aggregate history. This means agents react to changes in their environment, making decision making context dependent. These decision contexts are important all levels of modelling, whether in an abstract model, as the one discussed in this paper, or a more concrete model of e.g. migration in Klabunde (2014).

ABM also has increasing traction in choice modelling. Willekens et al. (2017) discuss four ABM of choices ranging from spouse selection to migration. They focus on the procedural aspect of decision making. By explicitly simulating agents over time, ABM can grapple with the procedural nature of choices. Klabunde and Willekens (2016) discuss four choices relevant to ABM of migration - equally relevant to most other data driven ABM: a) how expectations are formed, b) how choices are evaluated, c) how complex the decision is and d) whether networks are involved. This set of questions points to the huge variety of levels of abstraction for ABM. Dynamics or processes can be understood at different levels of generality. Bruch and Atwell (2015) discuss this complexity as *high* and *low dimensionality*. They distinguish abstract, simple and high dimensional worlds in modelling. For example the Schelling model of Segregation (Schelling, 1971) is an abstract model with a uniform choice setup in which agents simply look at the composition of their neighbourhoods and decide whether to move or not. The complexity of the model is entirely a result of the spatial interactions. Abstract or low dimensional models have little empirical realism, and their purpose is to “develop new ways of thinking about a problem and provide a great deal of theoretical stimulation for existing empirical research” (Bruch and Atwell, 2015, 6). On the other hand high dimensional models need high empirical realism, in particular on the behaviours of agents. And this is where Bruch and Atwell (2015) point to a well known problem of ABM: although social scientists have a lot of empirical data, the data and knowledge are rarely of the right kind for the validation of an agent-based model (Bruch and Atwell, 2015, p.7). The authors point out further that the validation of models of different dimensionality is rather different. Models of low dimension should not pretend to recreate reality but evaluate the relationships and mechanisms the model is based on (Bruch and Atwell, 2015, p.17). Rather they should be seen as models exploring the consequences of an assumed micro behaviour. This is where the model in this paper fits in, exploring the consequences of the assumption of (conspicuous) team reasoning.

Simulation models in general, and ABM in particular, have long been used to investigate the viability of cooperation (Axelrod, 2000a; Hoffmann, 2000; Gotts et al., 2003), and more specifically the conditions for the emergence or evolution of cooperation in social dilemmas (Axelrod and Hamilton, 1981; Boyd and Richerson, 1989; Skyrms, 1996, 2000; Nowak and Sigmund, 1998; Rand and Nowak, 2011).

Axelrod (1984) early computer simulations are often seen as the starting point of computational cooperation research. Different strategies are set against each other in an iterated prisoner’s dilemma to see which strategies survive. This is developed as a semi-evolutionary approach for the emergence of cooperation as it does not have mutation, meaning that current strategies cannot be invaded by new strategies or strategies that have died out. Since then plenty of computational work followed on the dynamics of cooperation based on these early computer simulations (see Axelrod, 2000b) and there are plenty of refinements and extensions (see below).

Evolutionary game theory, such as Skyrms (2004), uses replicator dynamics to test the survival of cooperative strategies. Boyd and Richerson (1989) show that cooperation does not evolve as an evolutionary strong strategy without any signalling and player selection. Nowak and Sigmund (1998) show that, with image scores, cooperation is a solid strategy to evolve in a two-player donor game. Large populations with lots of interactions are particularly favourable conditions for the emergence of cooperation. Other genres investigate inter-social influence on agent choices, such as punishing (e.g. Helbing et al., 2010) and sanctioning (e.g. De Pinninck et al., 2008), or prosocial preferences such as altruism (e.g. Németh and Takács, 2007) or fairness (e.g. Schindler, 2012).

Spatial game theory investigates environmental factors on cooperation dynamics, placing the agents into interaction spatial neighbourhoods. Flache and Hegselmann (2001) look at spatial aspects, analysing the influence of grid structures on the survival of cooperative behaviour. Hegselmann and Will (2012) model emergence and survival of cooperation using dynamic social networks. Environmental influences can also be investigated by placing the agents into particular interaction neighbourhoods (e.g. Nowak et al., 1994).

A third genre investigates inter-social influences on agent choices, such as the “shadow of the future” in iterated games (Axelrod, 1984), signaling of trustworthiness (Macy and Skvoretz, 1998) and agent reputation supporting cooperation (Raub and Weesie, 1990; Castelfranchi et al., 1998; Hales, 2002). The last two approaches simulate an inter-agent dynamic via more or less accurate type detection as a foundation of conditional cooperation. The role of reputation has been explored with more elaborate communication strategies, exploring for example the role of gossip (e.g. Giardini et al., 2014). Reputation can be slow to establish, especially in larger societies, and is potentially computationally costly, making gossip an efficient way of speeding up the process. Giardini et al. (2014) analyse gossip strategies at different group sizes, namely 5, 10 and 25. Groups of 10 members perform worst in their model. They hypothesise that this is because there are too many members for direct

reciprocity and too few to make the gossip mechanism relevant. Reputation and signalling strategies presuppose identifiability of individual free-riding agents.

An approach going away from the exclusive use of payoffs is the indirect evolutionary approach developed in Güth et al. (2000) and Güth and Kliemt (2000), motivations exist in parallel to the payoffs. A model for the evolution of trust is developed based on the Trust Game. In addition to the evaluation of payoffs, a motivational component m is introduced, a direct incorporation of a moral dimension into the preferences of the agents. Using a threshold on m , the population is divided into trustworthy and non-trustworthy types and an agent's decision is made conditional on what type it interacts with. Rather than using direct replicator dynamics on strategies, strategies become more frequent by agents actively changing their preferences, a form of learning leading to a two level dynamic, one on the level of strategies, one on the level of preferences.

These genres, evolutionary, spatial and inter-social, have been combined in numerous ways, in particular the influence of spatial settings on evolutionary or inter-social dynamics (e.g. Killingback and Doebeli, 1996; Skyrms, 2004). What they have in common is that they investigate the influence of the respective setting on the decisions of agents directly, i.e. the motivation to cooperate or defect is directly read-off the payoffs.

The model discussed in this paper is a contribution to ABM social dilemma games. It does go away from the tradition of iterated games, focussing on using payoff considerations as the driver of behaviour. Operationalising the idea of team reasoning, it models an agent's focus on being part of a team or not. The driving dynamics are about behaviour convergence, making it a part of the general literature of convergence models, e.g. opinion dynamics.³ Whilst the dynamics driving the model are social influence dynamics, in addition to macro behaviours such as length of time until a simulation converges etc, the outputs measured are also about levels of individual wealth and social inequality, outputs measured in normal iterated games.

4 The Model

Taking the idea of directly incorporating a collective dimension into agents' decision-making, the model starts from the assumption that agents consider themselves part of a team. It furthermore assumes that being part of a team increases an agent's likelihood of cooperation. It finally assumes that people observe how other members of the team behave, adapting their probability of cooperation to levels of cooperation found in the group.⁴

³For a detailed discussion of the relationship of the formalism of OD and TR, see the Appendix.

⁴We need to stress at this point that the model does not implement Sugden's team reasoning, as the model is explicitly designed to grapple with the possible anonymity of commons dilemmas. It implements a form of circumspect team reasoning akin to Bacharach's

In the model agents play a standard PGG where, in each round, each agent makes the decision to contribute to the public good or not. The contributions are multiplied by some factor (in our case, 3) and then shared equally between all agents. This means that free-riding agents get the same payout as those contributing, *but with paying in less*, increasing their individual payoff. If too many agents free-ride, there is no public good to share out and the game ends. Note that the implementation of the PGG in this model is done as the limiting case also known as an n -person Prisoner’s Dilemma where agents either decide to cooperate (contribute 1) or defect (contribute 0) rather than deciding what amount of their wealth they are contributing to the public good. The rationale behind this is to start off with a discrete choice model which allows to explore the most extreme gains and losses of team reasoners and individualists in the question about viability of team reasoning.

Agents are initialised as either *team reasoners*, with an initial cooperation probability of 1, and *individualists* with an initial cooperation probability of 0. From the size of the payout received, an agent can infer the cooperation level of the team, i.e. how many other agents have contributed and how many are free-riders. The level of free-riding influences the agent’s probability of their cooperation in the next round. There are no utility calculations directly influencing an agent’s choice. The payoff is simply used as a proxy for the cooperation levels within the team.

The model is written in NetLogo (Wilensky, 1999) and is available on Open-ABM (once the paper is published). The simulated process is straightforward. Each step of the simulation represents one round of playing a public goods game. A round happens in three stages.

1. Each agent decides whether to cooperate or defect, according to its own probability of cooperation;
2. Each agent gets the payout from the collective good;
3. Depending on the cooperation levels within the team and their personal speed of adaptation (**recency-bias**) each agent updates their probability of cooperation.

Players act in some arbitrarily fixed order, but since there is no temporal dependency within each of these three steps, they could just as well act in parallel or in random order. The **wealth** of each player (as they gain or lose wealth by playing the PGG) and their probability to cooperate change during the simulation. The cooperation probability mirrors an agent’s experience of other players playing as a team. Everyone starts with zero units of wealth. During the course of a simulation, the wealth of a player is allowed to go below zero. Players can differ in three attributes: their **wealth**, their **recency-bias** and their cooperation-probability (Table 1).

definition, with agents advancing trust in the first step by simply cooperating, adapting to the collective behaviour of the group over time. The model shows the individual and team consequences of different team constitutions and adaptability rates.

Attribute	Domain
wealth	\mathbb{R}
cooperation-probability	$[0, 1]$
recency-bias	$[0, 1]$

Table 1: Player attributes

Agents are initialised with a **cooperation-probability** of either 0 (*individualists*) or 1 (*team reasoners*). As the game unfolds agents update their **cooperation-probability** depending on whether or not other agents are cooperating, i.e., working as a team. Defection incrementally reduces an agent’s conception of the game being collective, cooperation incrementally increases it.

The first step of each round is for agents to decide whether to cooperate or defect. Formally, the **cooperation-probability** P for player a on any given round is given by

$$P(a) = \frac{1}{|A| - 1} \cdot \sum_{\{x \in A | x \neq a\}} p(a, x), \quad (1)$$

where A is the set of players constituting the team and $p : A \times A \rightarrow [0, 1]$ is a function expressing the cooperation-probability as a relationship between players. A cooperating player contributes one token of wealth towards the “pot.” That token gets subtracted from their current wealth. A defecting player does not contribute anything. An agent’s wealth has no influence on its choice to cooperate or defect.

In the second step of the round, agents get rewarded. Let m be the multiplication factor and $c : A \rightarrow \{0, 1\}$ a function giving the contribution of a player on that round. The reward (r) for each player is simply the sum of all contributions, multiplied by m and divided by the number of players:

$$r = \frac{m}{|A|} \cdot \sum_{x \in A} c(x). \quad (2)$$

In the third and last step of the round, agents update their cooperation-probability. This probability can be updated at different speeds, as determined by the **recency-bias** of the players. The **recency-bias** represents the weight of the experience of the most recent round relative to the rounds played previously. A **recency-bias** of 1 means that only the most recent round is taken into account. The smaller the **recency-bias**, the more relative importance is given to steps further in the past. A **recency-bias** of 0 means no new information is taken into account.

Let α_a be the **recency-bias** of player a . The updated **cooperation-probability** p_{t+1} of player a towards other members of team A is given by:

$$p_{t+1}(a, b) = (1 - \alpha_a) \cdot p_t(a, b) + \frac{\alpha_a}{|A| - 1} \cdot \sum_{\{x \in A | x \neq a\}} c(x). \quad (3)$$

This means that although the function p is a function between individuals, the cooperation-probability P is based on the level of cooperation observed in the team in the current round and the previous cooperation-probability with respect to an agent’s past experience of the team, in proportions determined by the **recency-bias**. We do not explore that possibility in the present paper, but keeping track of individual cooperation probabilities would allow agents to adapt to changing team composition over time.

The model described above is a convergence model similar to opinion dynamics models.⁵ Although the general dynamic is similar, there are some important differences between the two kinds of convergence. For more detailed information, please see Appendix 1.

The model implements a version of team reasoning akin to Bacharach’s *circumspect team reasoning*. For Bacharach, understanding a game as collective, agents advance a level of trust into their fellow team members to cooperate, operationalised by an initial cooperation probability of 1. There does, however, come a point where an agent can suddenly switch its perception away from team reasoning (Bacharach, 1999, p.119), operationalised by the adaptation of the cooperation probability to the cooperation probability of the actual team. The interesting aspect of the model is that agents do not know the cooperation probability of their team but can only observe the binary mapping of individual cooperation probabilities and infer the cooperation in the long run.

4.1 Experiments

Experiments were run in BehaviourSpace⁶ over the parameter space shown in Table 2. Final values were reported at step 500 (unless otherwise stated). The analysis was conducted in Python and R.

Input Parameter	Values	Output Parameter	Range
Recency bias	0.1, 0.2, ..., 1.0	Step of convergence	$\mathbb{N}_{>0}$
Proportion of team reasoners	0.1, 0.2, ..., 0.9	Cooperation prob.	$[0, 1]$
Multiplication factor	3	Wealth	\mathbb{R}
Number of ticks	500 (5, 10, 50)		

Table 2: Input and output parameters of experimentation.

Each player is initialised with a **wealth** of zero. Each player is assigned a value of one or zero for the initial **cooperation-probability**. The **proportion-of-team-reasoners** parameter tells us what part of the population should have an initial **cooperation-probability** = 1. All other players will have an initial **cooperation-probability** = 0. If **proportion-of-team-reasoners** is

⁵Thanks to an anonymous reviewer to pointing out this similarity.

⁶BehaviourSpace is a built-in experimentation tool in NetLogo.

0.5, for example, half the population will be team-reasoners and the other half will be individualists.

The value for a player’s **recency-bias** $\in [0, 1]$ is determined by the **global-recency-bias**. The final parameter of the model is the **multiplication-factor** that is used for calculating the payoff of the public goods game. In our case, the value is held constant at 3 as variation in the multiplication factor does not result in any changes to the model outcomes. The proportion of team reasoners was varied from 0.1 to 0.9 and the global recency bias from 0.1 to 1.0 (all by intervals of 0.1), for 90 possible combinations. For teams of 100 players, we ran each combination 50 times, for a total of 4500 runs. The main outputs analysed are levels of cooperation, levels of wealth for team reasoners, individualists, and the overall population and the number of runs that finish early.

5 Results and Analysis

This section discusses results of the implementation of team reasoning in a PGG. It shows that the model replicates the general features of a PGG and presents the particular individual and social outcomes of players in a team reasoning scenario focusing on levels of cooperation, individual and group wealth, and team equity.

5.1 Individual Results

First of all we look at the viability of team reasoning as a strategy in a social dilemma. Of course, PGG can be solved easily if a certain percentage of agents simply keep cooperating regardless of free-riding. A useful model of a PGG needs to show that agents do not remain suckers (at least not for too long). The theory of team reasoning is sometimes criticised for presupposing agents to be irrational (see Bacharach, 1999, p.143). Bacharach develops a version of team reasoning with agents displaying a kind of *benefit of the doubt* towards team members by recognising a situation as collective and trusting that others will do so, too. He calls it *circumspect* team reasoning, expressing agents weariness to remain suckers. The model implements this kind of *circumspect* team reasoning by stipulating some agents behave *as if* they were part of a team, but over time they adapt their behaviour at different speeds, depending on the recency bias.

Figure 1 shows the wealth differential for initial team reasoners for different group constitutions and recency biases. The figure shows that team reasoners should adapt to their environment fairly quickly as low recency-bias results in larger payoff differentials. It also shows however, that only low recency bias (≤ 0.3) lead to noticeable differentials over time.

As stated in the introduction, for a strategy to be viable, the player must not end up in a continuous sucker situation. Figure 2 shows changes to the cooperation probability of team reasoners, showing clearly how quickly players adapt to the social context depending on the recency bias. When team reasoners are in a majority they rarely go off the rails to defect. When team reasoners




Figure 1: Mean cooperation probability vs. wealth over time (steps 5, 50, 500). A darker red shade means that initial team reasoners do relatively worse than individualists.

are in a minority they adapt fairly quickly to the situation, minimising wealth differentials. As stipulated in Bacharach’s *circumspect* team reasoning, agents wearily advance cooperation but only remain cooperative if this behaviour is returned (see fig. 2); they do not remain suckers in adverse circumstances.

5.2 Societal Results

By definition of a social dilemma, more cooperation leads to more societal wealth. The hallmarks of collective dilemmas are that cooperation sustains the collective good, benefiting everyone, whilst defection benefits free-riding individuals as long as free-riding is relatively low and the public good persists. If free-riding takes over, the collective good perishes to the detriment of all.

Figure 3 shows the mean wealth accumulated over time for all agents. The darker shading shows that teams with high proportions of team reasoners do considerably better than teams with many individualists. The relationship is almost linear between wealth and levels of team reasoners in a team. Figure 3 shows that the model indeed replicates a social dilemma. It shows that the higher the number of team-reasoners, the higher the level of wealth and for settings with high initial probability of defection, levels of wealth are very low.

As the mean wealth is almost linearly related to the team composition, the recency bias seems to make next to no difference. This is different when we look at the median. Figure 4 shows mean and median levels of wealth for different levels of recency bias and different team compositions. The mean wealth is the same across all levels of recency bias. For slow and medium levels of change the mean and median are almost the same. For fast change the median deviates considerably from the mean, this is in particular the case for a recency bias of 1.⁷ This means that levels of wealth are heavily skewed for those values. For low proportions of team reasoners the skew is to low levels of wealth, for high proportions of team reasoners the skew is towards high levels of wealth. The reason for the skew is that for faster recency bias agent choices are more volatile, increasing the stochasticity of the model. For groups with equal numbers of team reasoners and individualists, these initial stochastic effects can lead to highly

⁷A recency bias of 1 is formally equivalent to the strategy commonly known as Tit-for-Tat (TFT).

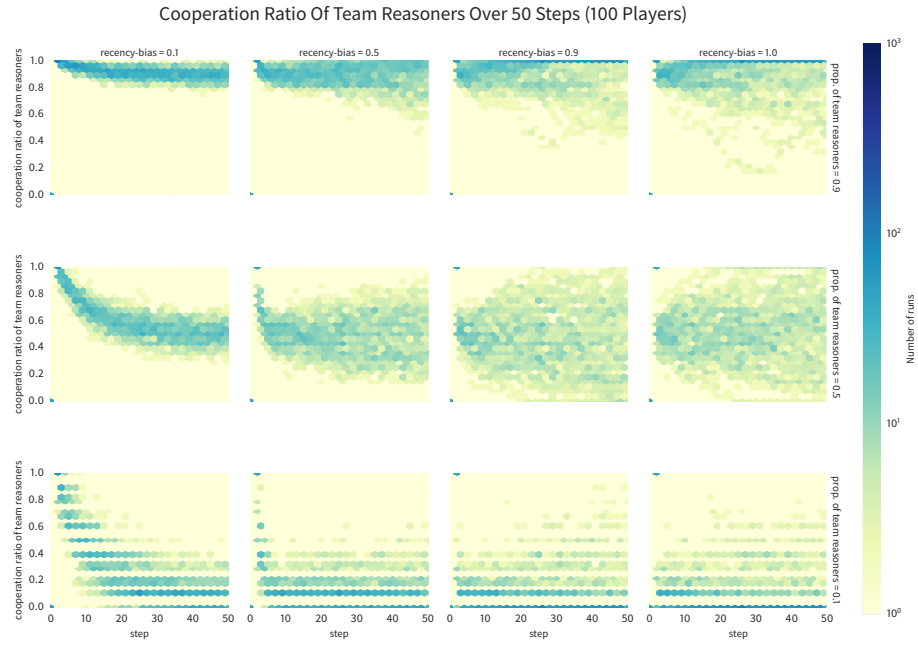


Figure 2: Cooperation Ratio of team reasoners over time for different team compositions.




Figure 3: Mean final wealth of players depending on team composition and recency bias.




Figure 4: Mean and median levels of wealth for different initialisations of the proportion of team reasoners and recency bias. Tit-for-tat is a recency bias of 1, fast = 0.9, medium = 0.5 and slow = 0.1.

skewed wealth distributions. The tipping point is at levels of team reasoning (> 0.5) and recency bias (> 0.5).

Both levels of team reasoning and recency bias thus have an effect on the equity of the team (cf. fig. 5). Figure 4 showed that increasing recency bias leads to more and more skewed distributions. Figure 6 looks deeper into these distributions by showing the Gini coefficient for different initialisations and its development over time.

Runs with a lower proportion of team reasoners have more inequality, and those inequalities are more persistent. The reason is that groups with high levels of team reasoning quickly end up in a fully cooperative equilibrium, meaning everyone has the same wealth. For groups with medium levels of team reasoning, inequality irons out over time, making the small differences from the first few steps irrelevant. For groups with low team reasoning, initial inequalities can persist because there is no wealth accumulation. Inequality is generally lower with more team reasoners (even at the beginning) because the payoff for everyone is bigger, making the differences resulting from having chipped in or not relatively smaller. These differences in Gini coefficient are made more salient by higher recency bias, since equilibria are reached more quickly. This explains the switch




Figure 5: Standard Deviation in Wealth. Note that the standard deviation is symmetric but the level of wealth is not. Whilst both extreme ends have equality, for low levels of team reasoners that is because there is no wealth, for high levels of team reasoners, because everybody is having the same share the full collective good.

from high Gini coefficient for low levels of team reasoners with low recency bias compared to high recency bias.




Figure 6: Mean Gini coefficient

Finally we look at the ultimate social outcome of loosing the collective good. Figure 7 shows the proportion of runs stopping early and the mean step of stopping when all agents in the group have payoff ≤ 0 .




Figure 7: Proportion of runs finishing early by proportion of team reasoners and recency bias.

It shows that only in high recency bias and low team reasoning environments does the public good cease. This shows that team reasoning, as implemented here, is a potential explanation for the high levels of cooperation found in experimental settings.

6 Conclusion and Future Work

This paper presents an agent-based model of team reasoning in a social dilemma. The model explores the interdependence of an agent’s individual choice and its social context. It maps a set of agent preferences to the binary choice of cooperation and defection. Repeated interaction leads to a convergence of the individual preferences and resulting behaviours. It contributes to choice modelling by showing pathways of using ABM to model heterogeneity, feedback, incomplete knowledge and emergence, here in the form of the emergence of cooperation. Heterogeneity of agents is represented by heterogeneous cooperation probabilities. Feedback between individuals and context results from the adaptation of an agent’s cooperation probability to the cooperation levels of their co-players. Incomplete knowledge is represented by mapping the preferences of agents to binary behaviour choices.

This team reasoning behaviour is an implementation of *conspicuous* team reasoning where agents can cooperate in one shot games, agents do not start defecting in situations of high cooperation and agents adapt to their team’s behaviour (cf. Bacharach, 1999, 2006). Players adapt their probability of cooperation to their experience of other players in their team. Team reasoners advance trust in their fellow players to cooperate, expressed in the model as an initial cooperation probability of 1. Given this advanced trust, team reasoners can be suckers, but as the behaviour is adaptive, the sucker status does not persist for long (fig. 2).

The model is theoretical, with all the well known limitations resulting from that (Sun et al., 2016). Nonetheless it contributes to the modelling of adaptation between stated and revealed preferences in situations of interdependence and social influence. The setup of a PGG is such that choices are strongly interdependent. Modelling individual’s adaptation of preferences through the interpretation of their co-players’ revealed preferences in repeated games allows to analyse the social influence on preferences in general. This is one of the fruitful areas for expanding the collaboration between choice modelling and ABM.

In addition to the theoretical implementation and exploration of team reasoning in an ABM, the model helps to investigate collective dilemmas in which anonymity plays a role. In anonymous collective situations social order mechanisms such as reputation or sanctioning of defectors does not work as defectors cannot be identified. Whilst large collections of people are usually not called groups or teams, it is the case that many collective dilemmas involve large numbers of people. Environmental examples like overfishing, the traditional grazing commons or emission targets are cases in point as are whip rounds in organisations where money is put anonymously into an envelope. The model is

a first implementation of team reasoning and is meant as a starting point for research in agent-based modelling using this new form of (collective) reasoning in agents, refocussing the modelling of cooperation on the social context rather than individual incentives and payoffs.

7 Acknowledgements

This paper has seen immense improvement through the comments of anonymous reviewers. We would like to thank in particular the reviewer who pointed out the similarity between opinion dynamics models and our model. It helped us to clarify some aspects and look for new ways of understanding the model. The research was funded by the AHRC/ESRC grant AH/L007282/1 “Collective Reasoning from a Moral Point of View”.

References

- Axelrod, R. (2000a). On Six Advances in Cooperation Theory. *Analyse & Kritik*, 22:130–151.
- Axelrod, R. (2000b). On six advances in cooperation theory. *Analyse & Kritik*, 22(1):130–151.
- Axelrod, R. and Hamilton, W. D. (1981). The evolution of cooperation. *Science*, 211:1390–96.
- Axelrod, R. M. (1984). *The Evolution of Cooperation*. Basic Books.
- Bacharach, M. (1999). Interactive team reasoning: A contribution to the theory of co-operation. *Research in Economics*, 53(2):117–147.
- Bacharach, M. (2006). *Beyond Individual Choice: Teams and Frames in Game Theory*. Princeton University Press.
- Badham, J., Chattoe-Brown, E., Gilbert, N., Chalabi, Z., Kee, F., and Hunter, R. F. (2018). Developing agent-based models of complex health behaviour. *Health & place*, 54:170–177.
- Balliet, D., Mulder, L. B., and Van Lange, P. A. (2011). Reward, punishment, and cooperation: a meta-analysis. *American Psychological Association*.
- Becker, G. S. (2017). *Economic theory*. Routledge.
- Bicchieri, C. (2006). *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge University Press.
- Boyd, R. and Richerson, P. J. (1989). The evolution of indirect reciprocity. *Social Networks*, 11:213–236.

- Brañas-Garza, P. and Espinosa, M. P. (2011). Unraveling public good games. *Games*, 2(4):434–451.
- Bruch, E. and Atwell, J. (2015). Agent-based models in empirical social research. *Sociological methods & research*, 44(2):186–221.
- Burton-Chellew, M., N., C. E. M., and West, S. A. (2016). Conditional cooperation and confusion in public-goods experiments. *Proceedings of the National Academy of Sciences*, 113(5).
- Burton-Chellew, M. N. and West, S. A. (2013). Prosocial preferences do not explain human cooperation in public-goods games. *Proceedings of the National Academy of Sciences*, 110(1):216–221.
- Camerer, C. F., Ho, T.-H., and Chong, J.-K. (2004). A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, 119(3):861–898.
- Cardak, B. A. (2004). Education choice, endogenous growth and income distribution. *Economica*, 71(281):57–81.
- Castelfranchi, C., Conte, R., Paolucci, M., et al. (1998). Normative reputation and the costs of compliance. *Journal of Artificial Societies and Social Simulation*, 1(3):3.
- Colman, A. M. and Gold, N. (2018). Team reasoning: Solving the puzzle of coordination. *Psychonomic Bulletin & Review*, 25(5):1770–1783.
- Colman, A. M., Pulford, B. D., and Rose, J. (2008a). Collective rationality in interactive decisions: Evidence for team reasoning. *Acta psychologica*, 128(2):387–397.
- Colman, A. M., Pulford, B. D., and Rose, J. (2008b). Team reasoning and collective rationality: Piercing the veil of obviousness. *Acta psychologica*, 128(2):409–412.
- Conte, R., Andrighetto, G., and Campennl, M. (2014). *Minding norms: Mechanisms and dynamics of social order in agent societies*. Oxford University Press.
- De Pinninck, A. P., Sierra, C., and Schorlemmer, M. (2008). Distributed norm enforcement via ostracism. In *Coordination, Organizations, Institutions, and Norms in Agent Systems III*, pages 301–315. Springer.
- Dekkers, H. P., Bosker, R. J., and Driessen, G. W. (2000). Complex inequalities of educational opportunities. a large-scale longitudinal study on the relation between gender, social class, ethnicity and school success. *Educational research and evaluation*, 6(1):59–82.
- Easley, D., Kleinberg, J., et al. (2010). *Networks, crowds, and markets*, volume 8. Cambridge university press Cambridge.

- Eek, D. and Biel, A. (2003). The interplay between greed, efficiency, and fairness in public-goods dilemmas. *Social Justice Research*, 16(3):195–215.
- Elsenbroich, C. and Verhagen, H. (2016). The simplicity of complex agents: A Contextual Action Framework for Computational Agents. *Mind & Society*, 15(1):131–143.
- Fehr, E. and Gächter, S. (2000). Cooperation and Punishment in Public Goods Experiments. *The American Economic Review*, 90(4):980–994.
- Fischbacher, U., Gächter, S., and Fehr, E. (2001). Are people conditionally cooperative? evidence from a public goods experiment. *Economics Letters*, 71(3):397–404.
- Flache, A. and Hegselmann, R. (2001). Do irregular grids make a difference? relaxing the spatial regularity assumption in cellular models of social dynamics. *Journal of Artificial Societies and Social Simulation*, 4(4).
- Gambetta, D. (2000). Mafia: the price of distrust. *Trust: Making and breaking cooperative relations*, 10:158–175.
- Giardini, F., Paolucci, M., Adamatti, D., and Conte, R. (2014). Group Size and Gossip Strategies: An ABM Tool for Investigating Reputation-Based Cooperation. In *Multi-Agent-Based Simulation XV*, pages 104–118. Springer, Cham.
- Gilbert, N. (2008). *Agent-based models*. Number 153. Sage.
- Gintis, H. and Bowles, S. A. (2011). *Cooperative Species: Human Reciprocity and Its Evolution*. Princeton University Press.
- Gotts, N. M., Polhill, J. G., and Law, A. N. R. (2003). Agent-based simulation in the study of social dilemmas. *Artificial Intelligence Review*, 19(1):3–92.
- Güth, W. and Kliemt, H. (2000). Evolutionarily stable co-operative commitments. *Theory and Decision*, 49(3):197–222.
- Güth, W., Kliemt, H., and Peleg, B. (2000). Co-evolution of preferences and information in simple games of trust. *German Economic Review*, 1(1):83–110.
- Hales, D. (2002). Group reputation supports beneficent norms. *Journal of Artificial Societies and Social Simulation*, 5(4).
- Ham, R., Junankar, P. N., and Wells, R. (2009). Occupational choice: Personality matters. Technical report, IZA discussion papers.
- Hardin, G. (1968). The Tragedy of the Commons. *Science*, 162(3859):1243–1248.
- Hegselmann, R. and Krause, U. (2002). Opinion dynamics and bounded confidence: models, analysis and simulation. 5.

- Hegselmann, R. and Will, O. (2012). From small groups to large societies: How to construct a simulator? Technical report, Working Paper, MODUS, Bayreuth.
- Helbing, D., Szolnoki, A., Perc, M., and Szabó, G. (2010). Punish, but not too hard: How costly punishment spreads in the spatial public goods game. *New Journal of Physics*, 12(8):083005.
- Heppenstall, A. J., Crooks, A. T., See, L. M., and Batty, M. (2011). *Agent-based models of geographical systems*. Springer Science & Business Media.
- Hoffmann, R. (2000). Twenty years on: The evolution of cooperation revisited. *Journal of Artificial Societies and Social Simulation*, 3(2):1390–1396.
- Hooff, J. A. R. A. M. et al. (2001). *Economics in nature: Social dilemmas, mate choice and biological markets*. Cambridge University Press.
- Janis, I. L. and Mann, L. (1977). *Decision Making: A Psychological Analysis of Conflict, Choice and Commitment*. New York: Free Press.
- Janssen, M. and Ostrom, E. (2006). Empirically Based, Agent-based models. *Ecology and Society*, 11(2).
- Kahneman, D. (2011). *Thinking, fast and slow*. Macmillan.
- Kassarjian, H. H. (1971). Personality and consumer behavior: A review. *Journal of marketing Research*, 8(4):409–418.
- Killingback, T. and Doebeli, M. (1996). Spatial evolutionary game theory: Hawks and doves revisited. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 263(1374):1135–1144.
- Klabunde, A. (2014). Computational economic modeling of migration. *Available at SSRN 2470525*.
- Klabunde, A. and Willekens, F. (2016). Decision-making in agent-based models of migration: state of the art and challenges. *European Journal of Population*, 32(1):73–97.
- Liu, L. and Eck, J., editors (2008). *Artificial Crime Analysis Systems: Using Computer Simulations and Geographic Information Systems*. Hershey, PA: Idea Group Publishing.
- Luebker, M. (2014). Income inequality, redistribution, and poverty: Contrasting rational choice and behavioral perspectives. *Review of Income and Wealth*, 60(1):133–154.
- Macy, M. W. and Skvoretz, J. (1998). The evolution of trust and cooperation between strangers: A computational model. *American Sociological Review*, pages 638–660.

- Németh, A. and Takács, K. (2007). The evolution of altruism in spatially structured populations. *Journal of Artificial Societies and Social Simulation*, 10(3):4.
- Nowak, M., Bonhoeffer, S., and May, R. M. (1994). More spatial games. *International Journal of Bifurcation and Chaos*, 4(1):33–56.
- Nowak, M. A. and Sigmund, K. (1998). Evolution of indirect reciprocity by image scoring. *Nature*, 393:573–577.
- Ormerod, P. (2011). *Butterfly economics*.
- Ostrom, E. (1990). *Governing the commons: the evolution of institutions for collective action*. Cambridge New York: Cambridge University Press.
- Ostrom, E. (2014). Collective action and the evolution of social norms. *Journal of Natural Resources Policy Research*, 6(4):235–252.
- Poteete, A. R., Janssen, M., and Ostrom, E. (2010). *Working Together: Collective Action, the Commons, and Multiple Methods in Practice*. Princeton University Press.
- Rand, D. G. and Nowak, M. (2011). The evolution of anti-social punishment in optional public goods games. *Nature communications*, 2:p.434.
- Raub, W. and Weesie, J. (1990). Reputation and efficiency in social interactions: An example of network effects. *American Journal of Sociology*, 96(3):626–654.
- Schelling, T. C. (1971). Dynamic models of segregation. *Journal of mathematical sociology*, 1(2):143–186.
- Schindler, J. (2012). Rethinking the tragedy of the commons: The integration of socio-psychological dispositions. *Journal of Artificial Societies and Social Simulation*, 15(1):4.
- Skatova, A. and Ferguson, E. (2013). Individual differences in behavioural inhibition explain free riding in public good games when punishment is expected but not implemented. *Behavioral and Brain Functions*, 9(1):3.
- Skyrms, B. (1996). *Evolution of the social contract*. Cambridge University Press.
- Skyrms, B. (2000). Game theory, rationality and evolution of the social contract. *Journal of Consciousness Studies*, 7(1-2):269–284.
- Skyrms, B. (2004). *The stag hunt and the evolution of social structure*. Cambridge University Press.
- Squazzoni, F. (2010). The impact of agent-based models in the social sciences after 15 years of incursions. *History of Economic Ideas*, XVIII(2).
- Sugden, R. (2003). The Logic of Team Reasoning. *Philosophical Explorations*, 6(3):165–181.

- Sugden, R. (2011). Mutual advantage, conventions and team reasoning. *International Review of Economics*, 58(1):9–20.
- Sun, Z., Lorscheid, I., Millington, J. D., Lauf, S., Magliocca, N. R., Groeneveld, J., Balbi, S., Nolzen, H., Müller, B., Schulze, J., et al. (2016). Simple or complicated agent-based models? a complicated issue. *Environmental Modelling & Software*, 86:56–67.
- Troitzsch, K. G. (2018). Can lawlike rules emerge without the intervention of legislators? *Frontiers in Sociology*, 3:2.
- Watts, C. and Gilbert, N. (2013). *Models of Genuine Novelty*. Edward Elgar Publishing: Cheltenham.
- Wilensky, U. (1999). NetLogo.
- Willekens, F., Bijak, J., Klabunde, A., and Prskawetz, A. (2017). The science of choice: an introduction.
- Yang, L. and Gilbert, N. (2008). Getting away from numbers: Using qualitative observation for agent-based modeling. *Advances in complex systems*, 11(02):175–185.

Appendices

A Appendix

The model discussed in this paper is a model of convergence. Over time the cooperation probabilities of agents converge. This kind of dynamic has been investigated extensively in the area of opinion dynamics (Hegselmann and Krause, 2002). Although similar, the models diverge in some important ways.

A.1 Individual Runs

A generalised opinion dynamics model uses a set of agents with opinions and a set of weights for each agent towards another. In their description of the classic Opinion Dynamics (OD) model, Hegselmann and Krause (2002) use the matrix M_{OD} below:

$$M_{OD} = \begin{bmatrix} 1/2 & 1/2 & 0 \\ 0 & 2/3 & 1/3 \\ 1/4 & 0 & 3/4 \end{bmatrix}. \quad (4)$$

This matrix means that agent a_1 weights its own opinion as $1/2$ and agent a_2 's opinion as $1/2$ but agent a_3 's opinion not at all. Agent a_2 does not weight agent a_1 's opinion, its own as $2/3$ and agent a_3 's opinion as $1/3$, etc.

Let agents a_1, a_2 and a_3 have the following opinion profile:

$$OP_t = \begin{bmatrix} 0.8 \\ 0.4 \\ 0.3 \end{bmatrix}. \quad (5)$$

Using the matrix M_{OD} to calculate the weighted means starting from OP_t we can get the new opinion profile by simple matrix multiplication. The updated opinion profile is:

$$OP_{t+1} = \begin{bmatrix} 0.6 \\ 0.37 \\ 0.43 \end{bmatrix}. \quad (6)$$

Over time this matrix will converge to all agents having the same opinion which will lie somewhere between 0.8 and 0.3. An OD model can converge to any value in the interval $[0, 1]$.

Our model of Team Reasoning (TR) can be expressed in a similar matrix form, where α_i is the recency bias of agent a_i , and n is the number of agents.

$$M_{TR} = \begin{bmatrix} 1 - \alpha_1 & \frac{\alpha_1}{n-1} & \frac{\alpha_1}{n-1} & \cdots & \frac{\alpha_1}{n-1} \\ \frac{\alpha_2}{n-1} & 1 - \alpha_2 & \frac{\alpha_2}{n-1} & \cdots & \frac{\alpha_2}{n-1} \\ \frac{\alpha_3}{n-1} & \frac{\alpha_3}{n-1} & 1 - \alpha_3 & \cdots & \frac{\alpha_3}{n-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \frac{\alpha_n}{n-1} & \frac{\alpha_n}{n-1} & \frac{\alpha_n}{n-1} & \cdots & 1 - \alpha_n \end{bmatrix}. \quad (7)$$

For a recency bias of 0.6 and three agents we get:

$$M_{TR} = \begin{bmatrix} 0.4 & 0.3 & 0.3 \\ 0.3 & 0.4 & 0.3 \\ 0.3 & 0.3 & 0.4 \end{bmatrix}. \quad (8)$$

Let agents a_1, a_2 and a_3 have the following cooperation probability profile.

$$CP_t = \begin{bmatrix} 0.8 \\ 0.4 \\ 0.3 \end{bmatrix}. \quad (9)$$

Now further assume that the actual behaviour profile (BP) is a_1 cooperates, a_2 cooperates, and a_3 defects:

$$BP_t = \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}. \quad (10)$$

Now the matrix M_{TR} is used to calculate the weighted means using the CP_t for an agent's own cooperation probability and the mapping of BP_t for the cooperation probability of the other agents. The calculations for each agent are:

$$p_{t+1}(a_1) = 0.4 \times 0.8 + 0.3 \times 1 + 0.3 \times 0 = 0.62$$

$$p_{t+1}(a_2) = 0.3 \times 1 + 0.4 \times 0.4 + 0.3 \times 0 = 0.46$$

$$p_{t+1}(a_3) = 0.3 \times 1 + 0.3 \times 1 + 0.4 \times 0.3 = 0.72$$

The updated cooperation profile is:

$$CP_{t+1} = \begin{bmatrix} 0.62 \\ 0.46 \\ 0.72 \end{bmatrix}. \quad (11)$$

Assuming the behaviour profile is the same as before, the cooperation probabilities in the next round will be calculated as follows.

$$p_{t+2}(a_1) = 0.4 \times 0.62 + 0.3 \times 1 + 0.3 \times 0 = 0.55$$

$$p_{t+2}(a_2) = 0.3 \times 1 + 0.4 \times 0.46 + 0.3 \times 0 = 0.3$$

$$p_{t+2}(a_3) = 0.3 \times 1 + 0.3 \times 1 + 0.4 \times 0.72 = 0.89$$

The profile will thus not converge to a value between 0.8 and 0.3 due to the stochastic nature of the mapping of cooperation probabilities to binary behaviour choices 1 and 0. The TR model can only stabilise when converging to either 1 or 0. For all other values of the cooperation probability, the profile will still change.

A.2 Aggregated Runs

Despite this, it might be that the macro behaviour of the models is equivalent, i.e. that the average level of cooperation probability across a number of runs is the value of convergence to be expected from the OD model. This is indeed the case. The model is a special case of an OD model *on aggregate* where the proportion of runs converging to 1 or 0 is linearly related to the initial proportion of team reasoners.

The level of equivalence is, however, at such a high level of abstraction as to be a case of epiphenomenality, i.e. several micro behaviours can produce the same macro behaviour. Given that the questions the model tries to answer are not, as is the case, about the value of convergence, we think the equivalence is harmless. Questions explored by the TR model are about the influence of team composition and recency bias on levels of wealth and inequality, a question that is not asked in the OD literature.




Figure 8: Macro convergence of the TR model over 7357 runs with **recency-bias** equiprobably one of $\{0.1, 0.2, \dots, 0.9\}$. The mean cooperation probability over multiple runs is the same as the opinion value that a continuous OD model would be expected to converge to.