

STOCKHOLM SCHOOL OF ECONOMICS
Department of Economics
5350 Master's Thesis in Economics
Academic Year 2021–2022

Carrots, Sticks, and Other Tricks:

Understanding the effectiveness and robustness of social mechanisms for encouraging cooperation in a complex adaptive system

Thomas Atherton (41906) and Andreas Bergström (41884)

Abstract

Cooperation is critical for success in many economic interactions, but social dilemmas between the best outcomes for society and for individual decision-makers often arise. This thesis adds to existing game-theoretic and behavioural studies by evaluating cooperation in the public goods game (a canonical social dilemma) in a more complex, realistic system. Specifically, agent-based simulations are used to create a spatial grid of 2500 agents, who are given heterogeneous social preferences and learn adaptively from neighbours' performance, and the game is played for 1000 rounds. After calibrating the baseline simulation to mimic real experimental results, social mechanisms designed to improve average cooperation are tested, including punishment, reward, and reputation. Overall, three key findings are presented: first, reputation is the most reliably impactful mechanism for improving cooperation; second, medium-sized 'Goldilocks' policy parameters (e.g. fines, rewards) typically produce the best results; and third, adding economic features like the *Homo Moralis* utility function to this interdisciplinary research field produces much more realistic results. Finally, the simulation method is shown to be useful not only in detecting dynamically evolving patterns, but also in testing the external validity of results from traditional methods, highlighting the wide potential usefulness of this novel method in future research.

Keywords: Agent-based modelling, Behavioural, Complex networks, Cooperation, Evolutionary

JEL: C63, C71, C73

Supervisor:	Andreea Enache
Date Submitted:	14 May 2022
Date Examined:	30 May 2022
Discussants:	Annika Quester and Jiechao Sheng
Examiner:	Karl Wärneryd

Contents

1	Introduction	6
1.1	Research Question	7
1.2	Purpose & Motivation	8
1.3	Thesis Structure	10
2	Literature Review	11
2.1	Foundations of Cooperation in Economics	11
2.2	Behavioural Perspectives on Cooperation	12
2.3	Interdisciplinary Insights on Cooperation	15
2.4	Insights on Cooperation from Agent-Based Simulations	17
3	Methodology	25
3.1	Methodological Approach: Agent-Based Simulation	25
3.2	Conceptual Framework	29
3.2.1	Game Setup and Purpose	29
3.2.2	Agent Definition – Heterogeneity and Utility Functions	30
3.2.3	Agent Definition – Strategies and Updating Rules	32
3.2.4	Game Mechanisms	33
3.2.5	Game Procedure and Outcomes	35
4	Model	37
4.1	MATLAB Implementation	37
4.2	Data	37
4.2.1	Identifying Parameter Values from Previous Research	37
4.2.2	Calibrating Parameter Values from Previous Experiments	40
4.3	Game Procedure	42
5	Results	43

5.1	(Overly) Simple SPGG with <i>Homo Oeconomicus</i>	43
5.2	Baseline SPGG with <i>Homo Moralis</i>	44
5.3	SPGG with Punishment Mechanism	47
5.4	SPGG with Reward Mechanism	50
5.5	SPGG with Policing Mechanism	54
5.6	SPGG with Reputation Mechanism	57
5.7	SPGG with Moral Adjustment Mechanism	60
5.8	Robustness Tests	62
5.8.1	Extending Rounds to $T = 10000$	62
5.8.2	Varying Distribution of <i>Homo Moralis</i> Parameter	64
5.8.3	Varying Spatial Distribution of Agents at $t = 0$	66
5.8.4	Varying Size of the Public Goods Multiplier	67
5.8.5	Varying Parameter K in the Fermi Equation for Strategy Updating	68
5.8.6	Varying the Rate of Innovation in Strategy Updating	69
5.8.7	Adding Semi-Stochastic Dynamism to the Public Goods Multiplier	71
6	Discussion	74
6.1	Overall Trends	74
6.2	<i>Homo Moralis</i> and Bounded Rationality in the Spatial Public Goods Game	75
6.3	Social Mechanisms in the Spatial Public Goods Game	76
6.3.1	Reputation	76
6.3.2	Punishment and Reward	77
6.3.3	Policing	78
6.3.4	Moral Adjustment	79
6.4	The Public Goods Multiplier	79
6.5	Robustness in Simulating the SPGG	80
6.6	Limitations	80
7	Conclusion	82

7.1	Potential Applications	83
7.2	Areas for Future Research	84
8	References	86
A	Additional Results	97
A.1	Baseline Simulation Results with Fehr-Schmidt Inequity Aversion Utility Function	97
A.2	Baseline Simulation Results with Charness-Rabin Conditional Welfare Utility Function	98
B	MATLAB Program Files	99
B.1	Scripts	99
B.1.1	Baseline Simulation Model	99
B.1.2	Simulation with Punishment	99
B.1.3	Simulation with Reward	99
B.1.4	Simulation with Reputation	99
B.1.5	Simulation with Policing	99
B.1.6	Simulation with Moral Adjustment	100
B.2	Supporting Functions	100
B.2.1	Neighbours Function	100
B.2.2	Neighbours Max Function	100
B.2.3	Fermi Equation Function	100

Acknowledgements

In pursuing a novel thesis topic such as this, we expected a high degree of scepticism and pushback. The reality has proven to be very different. We are extremely grateful for the advice and support given to us by our supervisor, Assistant Professor Andreea Enache, whose encouragement and constructive feedback has shaped this thesis far beyond our original expectations. This sentiment extends much more widely across the Stockholm School of Economics's Department of Economics faculty, not least to Professor Anna Dreber, Associate Professor Mark Voorneveld, Professor Magnus Johannesson, Assistant Professor Fabio Blasutti, and Professor Rickard Sandberg, all of whom offered targeted and insightful guidance in shaping both our initial idea and later implementation.

Our thanks also reach beyond the walls of SSE: to Associate Professor Erik Mohlin of the Institute for Future Studies, who offered particularly useful insights into the simulation modelling method; to Ingvar Strid, economist at the modelling division of Sveriges Riksbank, for his careful review of our work and advice on structuring results and robustness tests; and Gustav Karreskog, Postdoc in the Department of Economics at Uppsala University, for his ideas on finessing our research question and selecting a methodological approach.

Finally, our acknowledgements also go to the faculty and student body at the Stockholm School of Economics more broadly, for stimulating the intellectual curiosity to pursue an interdisciplinary thesis topic that we hope will appeal across multiple departments.

1 Introduction

“Competition has been shown to be useful up to a certain point and no further, but cooperation, which is the thing we must strive for today, begins where competition left off.”

– Franklin D. Roosevelt, 1912

Cooperation is a favourite topic for political speeches and policy debates, and has long been argued by politicians, economists, and biologists alike to be necessary for the very survival and evolution of the human race. It is as true today as ever, yet the last 10-15 years have been less supportive of this imperative. Although cooperation is crucial for the sound functioning of a globalised economic system, the post- Global Financial Crisis era has seen countries reverting to protectionist, anti-globalisation trade policies. Despite a clear scientific and economic rationale for taking action against climate change, international conventions like the annual UN Conference of the Parties have persistently fallen short of expectations due to fractious negotiations between parties unwilling to cooperate. Ever-changing cooperation dynamics between oil-producing nations in setting oil production levels has recently led to both all-time highs and all-time lows in oil prices. Most recently, Europe has witnessed a grave threat to its geopolitical security from the Russian invasion of Ukraine, after a decades-long decline in international cooperation with Russia. These are just a few examples of how anti-cooperative social dilemmas can arise between economic actors, and they hint at how complex the underlying drivers of cooperation can be.

If cooperation in social dilemmas can be beneficial for all, why then do decision makers so frequently get stuck in these unfavourable outcomes? Economics has revealed that some failures in cooperation could arise from self-interested, short-termist, hyper-rational opponents in zero-sum games like the famous prisoner’s dilemma. However, the real world is VUCA (Volatile, Uncertain, Complex, and Ambiguous), making it hard to assess cooperative dynamics with these hyper-rational tools, and other disciplines offer different explanations: for example, psychologists have found a behavioural tendency for people to overvalue holding their current endowments more than the potential greater benefits they would receive from cooperative strategies (Kahneman *et al.*, 1990), while evolutionary biologists point out that it is harder for individuals to cooperate until cooperation is already widespread in a community (Bernard *et al.*, 2016). All of these explanations are potentially valid in the right context, so these interdisciplinary considerations of cooperation in a VUCA world could be highly relevant to economic decision-makers. Traditionally, though, the economic discipline has largely shunned insights from other disciplines, at least until the recent advent of behavioural economics and its notable popularisation by Kahneman (2011), and while fields of evolutionary game theory and complexity economics exist, they have yet to take significant account of these complex evolutionary dynamics emphasised elsewhere.

Nevertheless, as highlighted at the 2021 Nobel Symposium “100 Years of Game Theory”, there is now a need to explore new techniques after 100 years of more abstract, deductive reasoning has solved most of the challenges available for these analytical methods. Kandori (2021) poses a ‘fundamental difficulty’ of studying dynamic adjustment processes towards Nash equilibria in the context of human decision-making using analytical game theory, because unlike machine algorithms, human decisions can go beyond a prescribed decision-making process and take into account newly emerging

factors like changes to their opponent's strategy, which adds instability into the dynamic adjustment process of their own strategies. Kandori argues that this has created a dead end for researchers who do not accept this instability issue or who dismiss such concerns as *ad hoc* and incompatible with their analytical methods, and proposes that the solution to breaking this deadlock is to engage more in empirical research, to find evidence of such emerging instabilities and even to potentially answer the question of how these unstable dynamic adjustment processes develop. Potential solutions for economics via computer science and interdisciplinary insights have been around in some forms for a while now, and there has indeed been a shift towards integrating multidisciplinary considerations to explore the complexity of economic phenomena in an uncertain, ambiguous world through computational methods. As these research endeavours continue to formalise into a distinct field of computer-simulated behavioural games, new insights into emergent patterns and out-of-equilibrium dynamics that cannot be solved using traditional analytical methods are now proving possible.

1.1 Research Question

Given the relevance of cooperation and emerging capabilities to simulate complex behavioural interactions between heterogeneous agents, this thesis seeks to address the following research question:

How effective are different social mechanisms (e.g. punishment, reward, reputation) for encouraging cooperation in the spatial public goods game in a complex system – where agents are boundedly rational and learn adaptively from their surroundings – and how robust are they to an uncertain, dynamically evolving environment?

This question directs focus to simulation methods, which are required for producing a spatial mapping of the canonical public goods game that is popular in behavioural economics for assessing cooperation between agents. This public goods game offers a richer environment for exploring cooperation dynamics than the simpler prisoner's dilemma game, insofar that it allows for more potential strategies and greater variability of game conditions. In particular, the public goods game allows for exploration of mechanisms like punishment, reward, reputation, and more, while offering a rather broad relevance to topics as diverse as tax contributions and conflict strategy. It also presents a classic situation of social dilemma, whereby group cooperation is desirable as a whole, but individuals face personal benefits if they defect from cooperation (Janssen and Ahn, 2006). Thus, it is a flexible game that enables researchers to address more complex questions around how to improve cooperation than could be solved with analytical methods, which also offers potentially more relevant policy implications for resolving social dilemmas.

The purpose of this research question can be captured in three broader goals, of which only the first is directly related to the topic of cooperation:

- (i) To better understand how economic agents cooperate in a realistic, complex world, both with and without social mechanisms to assist cooperation in social dilemmas;
- (ii) To more clearly establish the economic content of agent-based simulations like the spatial public goods game, which to date have been predominantly researched by evolutionary biologists, complex systems theorists, and physicists;

- (iii) To more clearly establish the complementary role that agent-based simulations can play in supporting the external validity of analytical and experimental findings in economic research, as well as in detecting novel emergent patterns and trends that analytical and experimental methods might otherwise miss.

At this point it is critical to state that this method is exploratory in nature. It does not claim powers of strong causal inference or the deductive closure of algebraic approaches; instead, it offers inductive insights into emergent patterns and dynamic processes that remain fallible to the complex interactions and influences of heterogeneous agents, and are subject to change or even obsolescence over time as the system evolves (Arthur, 2021). Hence, strong causal conclusions to the research question are not to be expected in the way that other methods might attempt to provide. However, as will be motivated next, the flexibility of this method might be precisely what is needed for addressing a research question that directly targets the complex, adaptive dynamics of human behaviour that exist in the real world, and the latter two goals of this thesis are designed to address a meta-goal of methodological pluralism in economics.

1.2 Purpose & Motivation


Cooperation is of fundamental importance within human interaction, and it spans a wide variety of economic applications. At a micro level, cooperation determines how individuals interact with others in their decision-making processes, and how firms strategise within their markets. At a macro level, cooperation plays a role in the arrangement of multilateral agreements and international negotiations on critical issues as systemic as peace or climate change. It therefore makes sense that economic decision-makers should research cooperation to be acutely aware of how it might develop in different settings, and what can be done to encourage more cooperation when a situation requires it.

To this end, economists have increasingly employed the methodological tools of game theory to better understand how agents interact and cooperate with each other, resulting in some major findings. Nash's (1951) famous Nash Equilibrium, for example, revealed how homogeneous, rational agents in the prisoner's dilemma game ought to be pushed towards a non-cooperative equilibrium of defecting against each other, while more recent game-theoretic models explore cooperation-based outcomes of more specific (though still abstractly modelled) problems, such as conflicts in managing scarce water resources (Madani, 2010). Similarly, behavioural economists like Fehr and Gächter (2000b) have employed the method of human experiments to see how agents actually cooperate (or defect) in simple games under different conditions. Overall, the economic study of cooperation has proven fruitful, and today, insights from these methods are regularly extrapolated to more advanced applications.

Yet economic decision-making rarely occurs in such simple, controlled algebraic or experimental contexts. Rather, the social reality in which heterogeneous agents typically make economic decisions is a complex adaptive system (CAS) with dynamics driven by agents changing strategies as they learn through interactions with others, and these systems are much more nebulous to define or solve (Holland, 2006). As Cilliers (1998) explains, these complex systems display a wide range of characteristics contradictory to conventional economic models: interactions are so numerous that they cannot easily be defined by systems of differential equations; they are nonlinear, making it

difficult to capture their functional form; they are dynamic over time, further complicating their specification; interactions are often localised, lacking knowledge of the entire system; feedback loops can be highly disruptive to equilibrium dynamics, so that the system is often far from equilibrium; the system may be open, ontologically speaking, thereby lacking definable boundaries; and so on. This results in many system properties that are not amenable to deductively solvable analytical methods like game theory, from the irreducibility of emergent patterns in the macroeconomy to its microfoundations, to chaotic dynamics that never settle at an equilibrium (Turner and Baker, 2020). Some great economists recognised these characteristics long ago: Kenyes's (1936) notion of 'animal spirits' referred to the complex interactions of investor sentiments that drive financial markets and are impossible to predict routinely, while Schumpeter's (1942) 'creative destruction' puts unpredictable, emergent patterns of creative destruction at the heart of progress in a capitalist system. Today, however, there is a considerable avoidance of these concerns in the teaching of economic methods (Colander, 2000), as leading economic pedagogues fall back on Friedman's (1953) epistemological philosophy of 'instrumentalism', disregarding concerns around realism in the modelling process so long as economic phenomena can be predicted precisely (if not reliably). But consider the implementation of rational, *Homo Oeconomicus* decision-making in this thesis's main game of focus, the spatial public goods game (SPGG) in Figure 1:

Figure 1: Evolution of the SPGG with only *Homo Oeconomicus* agents over 1000 rounds (blue agents are cooperators, while red agents are defectors).



The details of the game will be explained later in this thesis, but briefly, red dots represent defecting agents in the public goods game; blue dots are cooperating agents; and the different lattice snapshots

are taken at various rounds of the game (round 1, round 3, ... , round 1000). Experimental findings reveal average cooperation levels of around 40% in the game (Zelmer, 2003), whereas Figure 1 clearly tends towards 0% in both the short and long term. Evidently, then, rational game theory struggles to explain even basic cooperative dynamics in a simple behavioural game.

This position will be expanded in more detail in Section 3, but suffice it to say that there exist many reasonable arguments for exploring alternative techniques in economics that provide more ‘realism’, and as the applications of traditional game-theoretic methods reach a point of exhaustion, future studies in economics might have to take up more computational methods regardless. This is not to say that traditional methods are useless; on the contrary, it is widely recognised that continued advances in abstract game-theoretic and macroeconomic models have made these methods increasingly valuable for policymakers. However, for a subject that has historically sought to borrow mathematical methods from the natural sciences, surprisingly little attention has been paid in economics to the methods of physicists or evolutionary biologists like statistical mechanics or computer simulations to date, despite their apparent appropriateness for economics’s complex systems (Chick, 1998). Thus, by focusing on agent-based simulations of a popular behavioural experiment, this thesis hopes to explore how agents make economic decisions around cooperation in a more complex, dynamic, and ultimately realistic context, providing complementary insights to other studies via methodological pluralism and extending the scope of potential applications of game theory and other traditional methods.

1.3 Thesis Structure

To address this research question, Section 2 begins with a literature review of cooperation studies in economics, including game theoretic approaches, innovations from behavioural economics, and recent developments in the field of agent-based computer simulations. Section 3 then details the methodological approach taken in this thesis, providing epistemological justifications for using simulations to study cooperation in behavioural games, and offering a high-level theoretical framework for the model. In Section 4, the model is explored in more detail, addressing both the implementation of the theoretical framework in MATLAB and specific data sources and methods used for identifying and calibrating model parameters. Section 5 provides results from each model specification, and highlights relevant findings such as how cooperation dynamics evolve over time in different model contexts. Section 6 discusses the implications of these findings, demonstrating that while the simulation method is exploratory and not causal in nature, it is nevertheless useful for detecting emergent patterns or other dynamics that could offer insights for achieving cooperation in real-world settings. Finally, Section 7 concludes, summarising findings in relation to the research question and offering possible implications for cooperation studies and avenues for future research.

2 Literature Review

Cooperation has been a prominent topic within the study of economics, so the literature is both rich in detail and varied in scope. A brief overview of important research on cooperation, particularly within behavioural economics, is provided to make clear the foundations of this field. Then, attention is turned to interdisciplinary insights, given the pervasiveness of cooperation in other sciences and relevant findings from these works. This motivates the use of agent-based simulations to study cooperation, and sets up a deeper review of recent developments in this method, which are of significant influence in shaping the method and model presented in Sections 3 and 4 respectively.

2.1 Foundations of Cooperation in Economics

Cooperation has been a feature of human interaction for perhaps all of civilisation, but it was arguably only when humans began to learn more from each other and undergo ‘cultural adaptation’ that cooperation began to truly differentiate the human species from others, and as moral systems developed, the mechanisms of cooperation also evolved (Boyd and Richerson, 2009). While international trade, economic alliances, and other cooperative dynamics have all existed for many thousands of years, it is interesting to see how cooperation was considered in classical political economy, which remains influential today and often with nuance that analytical methods fail to express. Hobbes (2017 [1651]) is famous for his depiction of human nature being inherently selfish, describing life as ‘nasty, brutish, and short’, and the work of Hobbes has been a foundation of political economy ever since. Machiavelli (2014 [1532]) shared a similar view of predominantly selfish, self-interested human nature, though highlighted various forms of cooperation as tactically useful for achieving one’s own ambitions. Cooperation is even a prominent feature at the beginning of the economics profession, with the purported forefather of economics, Adam Smith. His book *The Wealth of Nations* (1982 [1776]) features much explanation of behaviour like the rational, market-driven competition arising from the ‘invisible hand’ or the cooperation that can generate benefits from the division of labour. However, Smith was writing at a time when economics was not mathematically formalised – while there was some mathematics, such as in the workings of the mercantilists or the physiocrats, it would not popularise rational, analytical methods until the Marginalist Revolution in the late 19th century (Screpanti and Zamagni, 2005) – and the philosophical reflections in his *Theory of Moral Sentiments* (2010 [1759]) were closer to the more realistic behavioural studies of today than much of what was achieved in the interim period. In this book, Smith postulates that humans have an innate sense of empathy towards others, so can feel concern or happiness towards those they interact with to the extent that there may be an innate desire to cooperate. Thus, through Smith alone there are two seemingly contrasting notions of cooperation: the rational, calculated cooperation driven by pecuniary economic gain on one hand, and the emotive, incalculable cooperation driven by empathy and moral values on the other.

However, following the rise of more analytical methods in economics, a distinct new form of rational decision theory came to the fore: game theory. The advent of game theory brought cooperation into an economic context in which it has remained ever since. It was von Neumann and Morgenstern (2004 [1944]) who popularised the field, modelling economic decision-making as game-like scenarios, with agents rationally ascertaining which potential strategy will maximise their utility, assuming that any opposing decision-makers will also be choosing their strategy rationally based on the

same principles of utility maximisation. These games have proven to be extremely flexible and offer insights into how rational agents might cooperate, with applications in situations ranging from international trade (Conybeare, 1984) to oligopoly (Corts, 1998), and even areas like conflict resolution (Snyder, 1971). Game theory's contributions to the study of cooperation have been major, though this original form of analytical game theory has suffered from a heavy criticism, namely its lack of realism. Strict assumptions of rationality and utility maximisation have limited the extent to which game theory can assess real-world economic decision-making, as it struggles to account for how social norms and values, imperfect rationality, and other environment-specific factors can influence the nature of cooperation (Jervis, 1988). Fortunately, behavioural economics rose to meet this challenge.

2.2 Behavioural Perspectives on Cooperation

Economic studies of cooperation took on a more nuanced approach with the rise of behavioural experiments. One reason for this is that the games of behavioural economics rest on less restrictive assumptions than traditional game theory. For instance, the former's notion of bounded rationality arose as a response to the latter's assumption of perfect rationality or 'logical omniscience' (Wheeler, 2020). Given humans' inability to always make perfectly rational decisions due to uncertainty, incomplete information, complexity in the decision-making environment, or simply the cost and time of computing all required information, bounded rationality is a core feature of behavioural economics's more 'realistic' models (Simon, 1964; Hacking, 1967). Similarly, it is more readily accepted that agents lack perfect self-control, with Thaler and Shefrin's (1981) theory of 'hyperbolic discounting' demonstrating agents' bias towards present rewards over future potential gains. A third tenet of much of behavioural studies is the existence of social preferences that influence decision-making, like the perceived fairness of a greedy economic agent that can drive others to punish them even if it results in a personal cost (Kahneman *et al.*, 1986). Already, these alternative theoretical foundations provide a deeper perspective on cooperation in economics: agents may act differently based on the incomplete information or uncertainty they are exposed to; they may choose to take for themselves in the short-term without recognising longer-term benefits of cooperation due to hyperbolic discounting; and agents may reject positive endowments from others if deemed unfairly low. These findings run contrary to traditional game-theoretic expectations – a perfectly rational agent, for instance, ought to take even the smallest of positive endowments from another, no matter how unfair it could be perceived to be – and clearly warranted further exploration.

These explorations predominantly take the form of behavioural experiments with human participants. Game-theoretic scenarios are played, sometimes for multiple rounds, and the conditions of the game are considered meticulously to control, as much as is possible, potential interferences with agents' decision-making processes. The prisoner's dilemma is one such game that can be used to assess cooperative habits under various conditions, revealing that with incomplete information as to opponents' altruism there can be more cooperation than predicted via rational game theory (Kreps *et al.*, 1982; Andreoni and Miller, 1993); repeated games can also lead to 'tit-for-tat' strategies dominating, under which agents copy their opponent's prior decision (Axelrod and Hamilton, 1981); and the level of cooperation can be affected by the length of repeated games (Normann and Wallace, 2011). The dictator game, in which one agent is given an endowment and chooses how much to share with another agent, is similarly malleable: Hoffman *et al.* (1996) showed that greater anonymity of dictators results in lower offers to their subjects, although cooperation is still

much greater than predicted analytically with rational expectations, while Ben-Ner *et al.* (2004) discovered that repeated rounds with reversal of roles between partners leads to highly correlated returns, even though subjects' matches are anonymous.

Most pertinent to the research of this thesis is the public goods game, in which a group of agents are given an endowment each round and may contribute this to a collective public goods pot, which is multiplied by a factor greater than one and split equally across all players. This has demonstrated, through various experiments, that: there is often only a minority of agents who refuse to pay into the pot and free-ride on the benefits from other contributors, running counter to rational expectations of non-cooperation (Fischbacher *et al.*, 2001); mechanisms like punishment and reward systems can encourage agents to conditionally cooperate, to various degrees (Fehr and Gächter, 2000a); and non-monetary punishments like shame or guilt can produce more cooperation than monetary punishments, possibly because monetisation makes the fine more socially acceptable, like a price for not cooperating, or because fines engender negative reciprocity feedback loops (Masclet *et al.*, 2003; Gneezy and Rustichini, 2000; Fehr and Rockenbach, 2003). Critically, for motivating this thesis's study of encouraging cooperation in the real world, Rustagi *et al.* (2010) found that these laboratory findings from the public goods game carry over into field experiments, with the presence of more conditional cooperators in forest management groups increasing productivity substantially. Equally, monetary incentives to cooperate can 'crowd out' by tarnishing the social reputation achieved from doing a 'good' thing, and gender can blur this situation further, with experimental evidence suggesting that women are more likely to display this crowding-out effect (Titmuss, 1970; Mellström and Johannesson, 2008). These experiments are not impervious to error, with environmental or other biases possibly influencing experimental subjects' behaviours or human error or motivations leading to malpractice in reporting statistically significant results (Simmons *et al.*, 2011). Nevertheless, all methods have their limitations, and the public goods game and other behavioural experiments offer insights into cooperation that might better approximate the reality of how individuals behave in making economic decisions.

Once experiments had been conducted, adapted functional forms can be more accurately defined for less than perfectly rational economic agents. Famously, Kahneman and Tversky (1979) redefined traditional notions of utility as a relative concept of 'value' relative to a reference point in their Prospect Theory. Through their value function, expressed in Equation 1:

$$V = w(p_1)v(x_1) + w(p_2)v(x_2) \quad (1)$$

under which the value V of a prospect is determined by the value of outcomes x_i in states $i \in [1, 2]$, with due consideration of each state's probability weights $w(p_i)$. In this specification, the relativity of value creates heterogeneous values for different agents and could arguably lead to different cooperative outcomes. Similarly, 'Inequity aversion' – the finding that some agents have a distaste for creating too wide a wealth gap between each other in the two-player dictator game and thereby cooperate more than expected in game theory – was modelled by Fehr and Schmidt (1999) as:

$$U_i(x) = x_i - \alpha_i \max\{x_j - x_i, 0\} - \beta_i \max\{x_i - x_j, 0\}, \quad i \neq j \quad (2)$$

with the players i and j , x_i is the utility from player i 's endowment, $\alpha_i \max\{x_j - x_i, 0\}$ determines how much negative utility arises from scenarios in which player j earns more than player i , and

$\beta_i \max\{x_i - x_j, 0\}$ captures the negative utility arising from inequity aversion in scenarios with player i earning more than player j . This is a famous utility function for capturing the notion of other-regarding preferences like altruism, and has received much recognition in the literature. Alternatively, Charness and Rabin (2002) characterise other-regarding preferences as:

$$\begin{aligned} U_i(x_i, x_j) &= (1 - \sigma + \theta)x_i + (\sigma - \theta)x_j, & \text{if } x_i < x_j \\ U_i(x_i, x_j) &= (1 - \rho + \theta)x_i + (\rho - \theta)x_j, & \text{if } x_i > x_j \end{aligned} \quad (3)$$

which can include the assumptions of standard competitive preferences ($\sigma \leq \rho \leq 0$), the above-referenced inequity aversion ($\sigma < 0 < \rho < 1$), social welfare preferences ($1 \geq \rho \geq \sigma > 0$) with no reciprocity, and negative reciprocity ($\theta > 0$; otherwise, $\theta = 0$). All of these expressions demonstrate the many directions in which economic research into cooperation can be expanded into more realistic settings.

Camerer and Ho (1999) offer a more advanced learning model of ‘experience-weighted attraction’, which allows agents to learn the strategies of those around them, attach different preference weights to these strategies based on their success, and probabilistically update their own strategy based on these attractions. The model is driven by two variables, $N(t)$ and $a_i^j(t)$:

$$N(t) = \rho \cdot N(t-1) + 1, \quad t \geq 1 \quad (4)$$

$$A_i^j(t) = \frac{\phi \cdot N(t-1) \cdot A_i^j(t-1) + [\delta + (1 - \delta) \cdot I(s_i^j, s_i(t))]}{N(t)} \cdot \pi_i(s_i^j, s_{-i}(t)) \quad (5)$$

Equation 4 defines the amount of ‘observation-equivalents’ of prior experiences, where ρ is a discount factor weighting prior periods, and Equation 5 determines agent i ’s attraction towards strategy s_i^j after round t , driven by discount factor ϕ , a dummy function I that equals 1 if $s_i^j = s_i(t)$ and 0 otherwise, and parameter δ that weights the payoff available from unchosen strategies (and therefore also weights the payoff from chosen strategy by $(1 - \delta)$). Finally, an example of other-regarding preferences being incorporated into agents’ utility functions in public goods games is Andreoni’s (1990) general utility function in Equation 6:

$$U_i = U_i(x_i, g_i, G), \quad i = 1, \dots, n \quad (6)$$

where agent i ’s utility is a function of consuming their private endowments (x_i), giving to the public good (g_i), and the total amount contributed to the public good itself, G , with $G = \sum_{i=1}^n g_i$. Thus, it is possible to capture utility not just from consuming the private good, but also the ‘egoistic’ or ‘warm-glow’ altruism that arises from giving to the public good, as well as the ‘purely altruistic’ utility arising from the total size of the public good. This expression tempers traditional assumptions in the public goods game and helps to approximate experimental findings, and is a great example of how back-and-forth insights from traditional, game theory, behavioural theory, and experimental findings can generate more nuanced and realistic mathematical functions to reflect how humans cooperate in economic interactions.

Today, this process of experimental exploration and subsequent refinement of the functional specification of agents continues. Yet as stated previously, there are some limitations to behavioural experiments. Beyond the usual epistemological concerns of errors or accidental or purposeful ‘p-hacking’ (Gelman and Loken, 2013), there are simply the logistical limitations of time and money (Charness *et al.*, 2013). Human experiments cannot be run indefinitely, nor with overly large groups of subjects, and the complexity of the game must remain tractable for efficient management and processing of the game and its results. This arguably leaves open a useful role for computer simulations to test the validity of experimental findings in more complex setups, and to integrate refined mathematical expressions for agents and their environments that have been adapted both from economics and from beyond.

2.3 Interdisciplinary Insights on Cooperation

Keynes (recounted in Heilbroner, 2000) once recalled that, when discussing economics with mathematician and philosopher Bertrand Russell, the latter asserted that he had considered entering the field, but thought it was too easy. However, in a similar discussion, quantum physicist Max Planck told Keynes that while he also had considered working in economics, he thought it too difficult. Planck’s point was that while the mathematics pursued in 20th-century economics is not as analytically demanding as that in other sciences like physics, the dynamics of real economic interactions can be far more complex than those found in the physical world. This sentiment has been echoed by Nobel Prize-winning economist Herbert Simon (1978, p.351), who noted in the context of his own time that economics is riddled with complex dynamics that create “intolerable computational burdens”, and further argued about the importance of an interdisciplinary approach to such problems (Simon, 1962). And yet, even 21st-century economics achieves a degree of popular notoriety for its abstract analytical methods and unrealistic models, which seek to linearise complex relationships and axiomatise simple assumptions that are expected to represent entire economies. Perhaps, then, it is worth investigating other sciences beyond economics for insights on cooperation.

Psychology is one of the closest disciplines to economics that deals with cooperation. Many of its insights have been discussed already under the guise of behavioural economics, but psychology has much more to say beyond this. Behavioural economics borrows from psychology, but still broadly retains classical economic assumptions; principles of rationality and self-interest, while softened somewhat, remain theoretical building blocks of behavioural theories. Conflicting concepts like altruism, for example, are often justified as rational acts of self-interest, with ‘altruistic’ agents who cooperate in games doing so for future prosocial acts to be directed towards them (e.g. in Hoffman *et al.*, 1996). Yet within the field of psychology, there is more licence for individuals to be truly irrational, prosocial, or otherwise nuanced. Other-regarding feelings like empathy are relatively more accepted as emotion without self-interested undertones, and the influence of society is seen to produce feelings like shame, which can in turn motivate cooperation without clear utilitarian drivers (Boyd and Richerson, 2009). Behavioural economics has done much to integrate considerations like these, but psychologists have explored such topics in greater depth. Integrating these phenomena further into behavioural economics studies via agent-based simulations or other methods could therefore offer deeper insights into how people interact in economic situations and what stimulates cooperation.

Biology, especially evolutionary biology, is another area of fruitful learning for economists. In fact, it

is conceivable to call economics a subset of ‘sociobiology’, as the seminal economist Alfred Marshall argued steadfastly (Hirshliefer, 1978). Biology and economics share a tendency to consider humans as being driven by a ‘selfish gene’ (Dawkins, 2016), and many biological models of cooperation are built on selfish tendencies that can be readily adapted into economics. One system often referred to is ‘kin selection’, where agents possess a tendency to cooperate with other agents sharing commonalities like familial ties (Maynard Smith, 1964); the matter of cooperation being conditional on receiving benefits in return plays a large role in evolutionary biology (see Noë and Hammerstein (1994), who actually use economic concepts to describe this phenomenon in ‘biological markets’); and cooperation can even arise as an unintended byproduct of selfish behaviour (Hamilton, 1971). Sachs *et al.* (2004) argue that multiple cooperative systems like these are often at work simultaneously in the natural world, and as a result, more recent agent-based models of cooperation have explored scenarios with two or more mechanisms. In the context of economics, Enke (2019) shows that groups with strong kin selection mechanisms tend to develop robust moral values and rule-based systems to regulate behaviour and promote cooperation. Similarly, Apicella and Silk (2019) highlight the shared usefulness of behavioural games to probe cooperative dynamics in complex systems in both economics and evolutionary biology. Along these lines, Dasgupta (2009) notes that economic agents’ ease of access to a “cooperative ‘infrastructure’”, or tools for credible enforcement of cooperation, plays a large role in determining how much agents will work together in a behavioural game. Evolutionary biologists have tested many different infrastructure tools like punishment and reward systems, sanctions, and reciprocity (both direct and indirect, or reputation-based) which can facilitate conditional cooperation, or non-conditional mechanisms like policing (West *et al.*, 2007), and findings suggest that such mechanisms can be influential in steering cooperation, albeit not necessarily in an evolutionarily stable way (e.g. Boyd and Richerson, 1992; Gardner and West, 2004). As will be seen, these learnings have been built into agent-based behavioural simulations, and conversely, simulations offer a useful sandbox for exploring the usefulness of these mechanisms in more complex environments.

Furthermore, in researching such complex environments, economics can learn from yet another of the natural sciences: physics. Economics has been said to suffer from ‘physics envy’, which is one of the reasons why it developed such distinctly abstract analytical methods (Davidson, 1984). But these methods are much closer to 17th-century Newtonian mechanics than 20th- or 21st-century physics – von Neumann and Morgenstern even argued that they wanted game theory to become the ‘Newtonian science’ underlying economic decision-making (Poundstone, 2021) – and may not always be suitable for modelling the complex dynamics of economics (Nelson, 2015), so there is much from the more recent canon of physics that remains to be integrated into economics. For instance, rather seeking to linearise the complex dynamics of cooperation amongst heterogeneous agents, economists could take note of how physicists deal with complex nonlinear relationships in the natural world. Chaos theory one example of this: it recognises that while there is inherent randomness in the natural world, there can be relatively simple underlying drivers that develop into highly chaotic, yet understandable, emergent patterns (Gleick, 1997). One example of this is the ‘butterfly effect’, a theory suggesting that small movements like a butterfly flapping its wings could trigger a chain reaction of events that lead to a tornado on the other side of the planet (Lorenz, 1963). There is already a sub-field of ‘econophysics’ that uses the tools of statistical physics to detect emergent patterns in economic systems (Burda *et al.*, 2003), like the discovery of scaling laws in financial markets (Brock, 2000a) and fat-tail distributions in derivative prices (Mandelbrot, 1963). Econophysics is closely related to complexity economics, but has not progressed significantly into the mainstream economics discipline. Yet the field of Agent-based Computational Economics (ACE),

which places greater emphasis on agent-based simulations, has started to bridge this interdisciplinary gap in a meaningful way, and will be explored further in Section 2.4. Agent-based simulations are evidently a methodological commonality between physics and evolutionary biology, and with the interdisciplinary insights of game theory, behavioural economics, and psychology, rich simulations could be developed to explore complex economic phenomena like cooperation.

Overall, these insights are of much use to complexity economics, and specifically its sub-field of agent-based behavioural simulations, which has become multidisciplinary in nature. This harks back to the very beginnings of complexity economics, which arose through a series of discussions on the state of economics in the late 1980s at the Santa Fe Institute between economists, evolutionary biologists, physicists, and other academics (Anderson *et al.*, 1988). During these meetings, researchers from outside the economic profession criticised the subject's narrow focus on abstract mathematical methods, and proposed alternative tools for exploring economic phenomena that might offer a better epistemological fit with social reality. Simulations, for instance, are highlighted in a later Santa Fe discussion as useful tools as they allow for tracking emergent patterns over time that can arise out of agents' interactions within complex economic systems (Arthur *et al.*, 1997). Yet despite multiple rounds of discussions, and some growing pockets of complexity research within the economics profession (notably a large faculty at the Santa Fe Institute, as well as Princeton's Studies in Complexity and Oxford's Institute for New Economic Thinking), the use of agent-based behavioural simulations has largely been shunned to niche research groups in computational or evolutionary economics, or to interdisciplinary journals in physics, statistical mechanics, or evolutionary biology – until recently.

2.4 Insights on Cooperation from Agent-Based Simulations

The first use of agent-based simulations in economics is arguably Schelling's (1971) model of spatial segregation. In this model, he demonstrates how agents that take one of two colours and possess a preference to be next to similarly-coloured agents on a grid space can, over time, move around to form remarkably segregated groupings on the grid (see Figure 2). The model's simplicity, combined with its ability to reveal emergent patterns like the unintended consequences (in this case, extreme societal segregation) that can arise at the macro level even if no single agent has strict preferences for segregation, enables it to succeed where traditional analytical methods struggle to identify such trends. For this work amongst others, Schelling received the 2005 Nobel Prize "for having enhanced our understanding of conflict and cooperation through game-theoretic analysis". Critically, though, Schelling (2006, p.4) points out that "I've never used much mathematics", yet despite this his game-theoretic contributions through simulations and other methods have been helpful in identifying how economic agents engage in strategic decision-making. Hence, Schelling's model can be understood as the first attempt to lead game theory and the study of cooperation into computationally-oriented methods, rather than traditional algebraic methods.

It was not until the 1990s that the first significant contribution to agent-based simulations arose specifically addressing questions around the strategy of cooperation and other behavioural economics questions. Nowak and May (1992) are credited with this achievement, through their simulation of a spatial prisoner's dilemma game involving 40,000 agents distributed on a 200x200 spatial lattice. In this game, each agent either always cooperates or always defects, with no more complicated strategies, memories, learning processes, or punishment/reward mechanisms involved. Each agent plays the

Figure 2: An updated representation of Schelling’s spatial segregation model, where F is the fraction of agents’ tolerance for how many of their own colour they desire around themselves. Reproduced from Hatna and Benenson (2012).





Figure 3: Even a simple spatial prisoner’s dilemma simulation with small changes in initial conditions can reveal starkly contrasting, chaotic dynamics of cooperators (blue agents), defectors, (red), and those changing strategy from one to the other (yellow/green), such as scenarios **a** versus **b**. Reproduced from Nowak and May (1992, p.827).



game with each of its immediate neighbours, and at the end of the round that agent’s site updates to reflect the strategy of the highest-scoring agents within its immediate neighbourhood. The game can then be iterated many time to explore asymptotic patterns for different initial conditions, and interestingly, despite chaotic dynamics across rounds, common long-term proportions of agents can be identified. Figure 3 is an example of two parameterisations after 200 rounds, where blue agents are successful cooperators, red are successful defectors, and yellow and green are unsuccessful agents who switched from one strategy to the other after the last round. This shows just how complex patterns of cooperation can become, even in a very simple rules-based game. The use of computational simulation and discussion of complexity in the study of cooperation gained further popularity through publications like Axelrod’s (1984) *The Evolution of Cooperation* and subsequent (1998) *The*


Complexity of Cooperation, and in reviewing the latter, the renowned economist Fukuyama (1998, p.142) admitted that the incorporation of interdisciplinary methods for complex adaptive systems is “overdue”. Before long, even prominent economists of more traditional, analytical methods like Nobel prize-winner Sargent (1993) began to support exploring relaxed assumptions like bounded rationality, heterogeneous and adaptive agents, and evolutionary programming.

Thus, from the 1990s it became clear that developments in computing power and methods like agent-based simulations were finally capable of managing and evaluating Simon’s “intolerable computational burdens”, and simulations of complex interactions began to arise in greater force. Notably, Epstein and Axtell (1996) of the Brookings Institution created an agent-based simulation called the Sugarscape, through which they hoped to lay the foundations for a ‘*transdiscipline*’ of computational complexity studies across economics, evolutionary science, and more. The model consists of agents on a two-dimensional grid featuring a rich topology of food (sugar) distributions – some areas on the grid are plentiful in sugar, others sparse. Agents possess rules to move around, search for food, mate, die, and compete for resources. Interestingly, over time these agents can learn to group into tribes, migrate away from overpopulated resource-rich areas, and engage in violent conflict over resources, producing rich dynamics and complex social histories. This in turn produces fascinating insights into cooperation dynamics, like the emergence of complex tribal groupings that cooperate within groups to compete against other groups and the roles of factors like seasonality and resource distribution, and the Sugarscape model is still available today via an open-source website for researchers across various disciplines to use as a sandbox laboratory for exploring the evolutionary dynamics of an artificial society ¹. Figure 4 provides one illustration of how the Sugarscape environment can develop over time. Axtell (interviewed in Colander *et al.*, 2004, p.265) conceded that the Sugarscape model was only meant to be a simple example of the possibilities of agent-based modelling, but its popularity upon release has been influential in later agent-based simulations.

In the 2000s, with a creeping formalisation of the research field of ‘Agent-based Computational Economics’ (Tesfatsion, 2002), clearer guides to agent-based modelling for economics and business applications (Heppenstall *et al.*, 2012), and increasing interest in agent-based simulations for quantitative finance (driven by influential papers like Arthur *et al.*, 1996), the study of cooperation through simulations of behavioural games began to mature. Validation of ‘explorative’, inductive simulations from prominent economists like Axelrod (2006) and formalisation of simulation frameworks by computational economists (e.g. Marks, 2007) also helped to spur the field on. For instance, Kirchkamp and Nagel (2007) tested simple learning and cooperation strategies in the prisoner’s dilemma game, and found that while simple *imitation* strategies do not necessarily make a significant difference, when applying strategies like imitating the most successful neighbour in a spatial game context (in which agents are distributed evenly on a circle edge), imitation can have more influential effects as agents of certain types can group together. Bramoullé and Kranton (2007) found that the agent network structure can influence cooperative dynamics in various ways. Similarly, Helbing and Yu (2009) investigated *success-driven migration* in a more advanced spatial prisoner’s dilemma game with empty spaces on the grid that permit agents to test new neighbourhoods by playing practice rounds with other agents in that area, and then move to that area if the conditions suit them better. Combined with the learning protocol of imitating successful neighbours, this migration mechanism could generate cooperative clusters even when conditions were designed to vastly favour defection, as in the simulations of Figure 5. Here, red cells are defectors, blue are cooperators,

¹The Sugarscape model is available at <http://sugarscape.sourceforge.net/>.

Figure 4: A simulation in the Sugarscape model, showing how spatially segregated and culturally distinct (red and blue) agents can learn to cluster into their own groups around resource rich areas (yellow cells) over time, and even depose a rival tribe from the landscape. Reproduced from Epstein and Axtell (1996, p.75).



green are defectors who became cooperators in the last round, yellow are cooperators who became defectors, white are empty spaces, and ‘noise’ refers to a heightened amount of randomness in the system. After 200 iterations it became clear that while imitation and migration mechanisms by themselves struggled to spread cooperation, the mechanisms combined could induce dominance of cooperation even in these adverse conditions, suggesting that mobility and imitation can encourage robust defensive clusters of cooperation. Furthermore, a related spatial prisoner’s dilemma game by Traulsen *et al.* (2010) revealed that randomness, rather than threaten cooperation by breaking up competitive clusters, can actually lead to higher rates of competition than achieved by imitation mechanisms alone, and as such, the randomness inherent in social systems can by itself explain some of the cooperation that is seen in the real world. Findings like this led Helbing and Yu (2010) to declare that agent-based game simulations can offer significant insights into puzzles long held by economists.

Figure 5: Even in conditions favouring defection, cooperation (blue cells) can be induced over 200 iterations by combining imitation and migration mechanisms. Reproduced from Helbing and Yu (2010, p.3681).






Figure 6: Two separate simulations in the top and bottom rows of a spatial public goods game with punishment. Top: agents can be defectors (red), cooperators (blue), or punishing cooperators (green), and **a**, **b**, **c**, and **d** are of $t = 10, 40, 150$, and 1000 . Bottom: agents can defect, cooperate, or be punishing defectors (yellow), and snapshots are of $t = 20, 200, 1000$, and 10000 . Reproduced from Helbing *et al.* (2010, p.10).



Around 2010, notable agent-based simulations of spatial public goods games began to appear more consistently to address the issue of social dilemma, with a wide range of designs. Many focused on *punishment* and *reward* mechanisms: Helbing *et al.* (2010) built a spatial public goods game with a punishment mechanism, in which agents can defect or cooperate and can also choose to levy a punishment on defectors at a personal cost. As seen in Figure 6, punishing cooperators (green) can crowd out non-punishing cooperators (blue) as they are more effective at combatting defectors (red), while punishing defectors (yellow) can also create a stable equilibrium of agent types after many iterations, but only if they exist in a good balance with cooperators (otherwise the system becomes unstable and there is no clear equilibrium). Meanwhile, Barr and Tassier (2010) highlighted the usefulness for economists of exploring concepts like agent movement in spatial games, which are hard to assess in experiments but can be easily evaluated in simulations. Szolnoki and Perc (2010) turned attention to a spatial public goods game simulation with a rewards mechanism, whereby agents can defect, cooperate, or cooperate plus give a reward to other cooperators at a personal cost. Similar to the punishment research, they found that small to moderate rewards can stimulate cooperation, but high rewards (and therefore high costs to the rewarders) can actually create unstable cooperative dynamics as agents switch between strategies more frequently. Ye *et al.* (2011) expanded on punishment mechanisms by adding a sympathy component to the public goods game that effectively compensates agents a small amount for the cost of punishment, to reflect positive recognition from sympathetic members of the community who show respect, support, or admiration towards those who endure the cost of punishment to punish defectors. As in Figure 7, it was shown that including this sympathetic behaviour can create strikingly more stable systems of punishment, with the proportion of ‘punishers’ (who are also cooperators) vastly outstripping non-punishing cooperators and defectors after a number of rounds.

Figure 7: Punishing cooperators (green) can vastly outstrip non-punishing cooperators (blue) and defectors (red) in the public goods game if they receive a small compensation driven by community sympathy. Reproduced from Ye *et al.* (2011).



From 2013 onward, the literature begins to explore more advanced components of spatial public goods simulations in different directions: Wang (2013) considers how agents with different learning abilities adapt their decision-making to imitate the collective performance level in a spatial prisoner's dilemma game, and discovers that a medium learning ability can enhance cooperation; Wardil and da Silva (2013) investigate how agents playing two different games in the same simulation can influence cooperation in each game; Wu *et al.* (2014) add an *influence* component to imitation strategies in spatial public goods games, whereby agents who are imitated a lot increase in influence, and therefore are imitated even more, and this positive feedback loop is found to be very supportive of cooperation; Lucas *et al.* (2014) simulate a public goods game based on agent types and parameterisations from


prominent behavioural economics experiments, and find that heterogeneity in social preferences can negatively impact cooperation; and Hagmann and Tassier (2014) show that allowing agent movement vastly improves the likelihood of them finding Pareto dominant Nash equilibria. In a rather novel and currently topical twist, Zhang (2013) even adapts a simulation to assess other-regarding preferences in a ‘spatial vaccination game’, while Waltman *et al.* (2013) utilise the simulation method to model cooperative behaviour amongst spatially distributed firms based on the level of information they have, demonstrating the vast flexibility of this method.

Then, from 2015, further advancements and refinements begin. For example, Wang and Chen (2015) simulate a reputational public goods game with two types of agents, those who spread their heuristic rules and those who follow to those rules. Chen *et al.* (2015) explore the spatial public goods game on a very sparse lattice, to understand how different movement rules can alter the outcome of aforementioned migration strategies. Chen *et al.* (2016) refine the influence mechanism with an *adaptive reputation* assortment, showing that higher reputational requirements for switching strategies promotes cooperation as this allows for more heterogeneous strategies and less herd mentality around defection. Evidence continued to mount on prior findings, such as with Cardinot *et al.* (2018), who argue that while heterogeneity by itself is not enough to encourage cooperation, if it increases the amount of overlap between similar agents it can indeed support cooperation. Building on punishment mechanisms, Zhang *et al.* (2018), recognising that agents are willing to tolerate some defection before punishing in a spatial public goods game, investigate this level of tolerance and find that variable *tolerance-based punishment* strategies can stabilise periods of high or full cooperation. On a similar note, Liu *et al.* (2018) proposed a *synergy punishment mechanism* through which punishment cost is reduced if many punishers collaborate, reflecting benefits of wider commitment to punishment like reduced social stigma and economies of scale. New ideas have also been introduced in the recent literature, like how Realpe-Gómez *et al.* (2018) combined Camerer and Ho’s (1999) experience-weighted attraction model of learning with additional considerations for social norms, *i.e.* what the socially appropriate decision to take is. They find that this combination of behavioural theories closely approximates real cooperative behaviour, and moreover is ambivalent to the interaction environment, be it a static square lattice or otherwise.

Finally, in the past couple of years there has been a trend of combining multiple different strategy updating rules and cooperation mechanisms into even more advanced simulations. Zhu *et al.* (2020a) build on issues of heterogeneous strategies by adding a reputation mechanism to a game environment of both imitation and *innovation* strategies (the latter meaning that agents randomly choose to switch their strategy irrespective of their surroundings), and reputation is found to resolve conflict and encourage cooperation. Meanwhile, Quan *et al.* (2020) combine the tolerance concept with the reputation mechanism, again finding that reputation-based imitation promotes cooperation and tolerance considerations further enhance it. He *et al.* (2020) also develop a reputation mechanism of agents achieving 1 point of reputation for every round of cooperation and losing 1 point for defecting, by incorporating a factor α representing *strategy persistence*, which effectively scales up the reputation component if agents persistently play the same strategy. They reveal in Figure 8 that a small amplifying value for the α parameter can promote sustained full cooperation, while larger values can force down cooperation after a periodic boost, albeit still reaching a higher proportion of sustained cooperation than without the persistence factor at all. Uniting behavioural theories with current-day simulations, Herings *et al.* (2021) allow agents to play strategies of altruism, egoism, and partial cooperation, and find that altruists can be effective enablers of partial cooperators. Overall, some trends in current simulations are clear: there is a movement towards combining multiple

strategies and mechanisms into one game; these combinations often increase cooperation, due to the multi-faceted incentives to cooperate that they offer; and medium-sized punishments, rewards, and reputation effects are often better for encouraging cooperation than larger ones, running counter to traditional theories of monotonically increasing returns from such mechanisms and demonstrating the usefulness of simulations in capturing the complex, dynamic patterns that can emerge.

Figure 8: *The proportion of cooperators (ρ_c) in a public goods game can be sustained at 100% within 100 rounds or Monte Carlo steps (MCS), if persistence-based reputation mechanisms are used with a modest persistence factor (α) of 2 (blue line). Higher cooperation levels are still maintained at greater values of α (yellow and green lines) than without it at all (red line). Reproduced from He et al. (2020, p.3).*



In retrospect, it is surprising that despite the support of complex adaptive systems, interdisciplinary insights, and simulation methods from multiple economists referenced here – some of whom reached the highest echelons of recognition as Nobel Laureates – the field of complexity economics and its computational simulations have remained on the outskirts of conventional economics until the past decade. This might have arisen due to the lack of requisite computational power until recently; many other, more critical explanations, also exist. Regardless, the subset of simulations reviewed here demonstrates how many possible strategies and environmental mechanisms can be utilised to promote cooperation, and it is clear from the literature today that there is ample justification for using simulations to better understand the dynamics of cooperation. Thus, while this method is still gaining popularity and has become considerably more refined in the last few years, there remains plenty of scope for exploring cooperation, as will follow in this thesis.

3 Methodology

This section focuses on high-level factors influencing the methodology underlying this thesis's model. A general overview of the agent-based computational economics and complexity economics research fields and the relevance of their inductive reasoning methods (including agent-based simulations) for modelling a real-world VUCA (Volatile, Uncertain, Complex, Ambiguous) environment is given in Section 3.1. Section 3.2 then sketches out the theoretical framework used in this thesis, linking the study of cooperation in behavioural economics more specifically to present-day research on agent-based simulations in economics and explaining choices behind the model's design.

3.1 Methodological Approach: Agent-Based Simulation

Due to the complex and dynamic nature of behavioural phenomena like cooperation, it strikes most outside the economics profession that assumptions like perfect rationality and static preferences are not entirely satisfactory, so a thorough understanding of cooperative dynamics in complex, real-world environments has not yet been reached. Moreover, increasing awareness that the economy is a highly uncertain, volatile system has raised interest in how to generate reliable policy insights in a VUCA context. Recognising these concerns alongside the clearly documented limitations of more traditional methods like behavioural experiments and game theory, this paper uses the method of agent-based simulations to explore the complexities of group cooperation dynamics. Yet this is not a common method employed by economists outside the realms of complexity economics, computational economics, or evolutionary economics (Colander, 2000). In fact, many articles like this are confined to niche journals of physics or statistical mechanics, or to broader scientific journals like *Nature* and *Science* (cf. Helbing *et al.*, 2010; Cardinot *et al.*, 2018; Arthur, 2021). Thus, some epistemological reasoning is required to justify why the simulation method is particularly appropriate for addressing the economics of cooperation.

First is the broad philosophical matter of *ontology*. Ontology is a subset of metaphysics, and covers the 'big' questions of what exists, how reality is constructed, and how things interact and relate to each other (Hofweber, 2021). While economics originated from the broader studies of philosophy and political economy, these sorts of considerations are often disregarded in economics today, with a general appeal to Friedman's 'instrumentalist' philosophy that waves away reality, so long as assumptions in models enable precise measurement of something. This is often still enough to obtain useful insights from economic research, but when dealing with complex systems, the reality of their complexity has to be recognised. In such instances, Davidson (1984, p.574) argued that he would rather be "roughly right than precisely wrong". A prominent Cambridge economic philosopher, Lawson (2003), has defined an economic ontology of 'critical realism' as systems characterised by 'openness', that is, systems in which dynamics can be influenced by undetectable disturbances from which unpredictable dynamics emerge. This runs counter to the 'closed' systems of deductive analytical models which, ontologically speaking, are defined by a limited number of axioms and variables which can only interact in fixed ways, such that no new influences or emergent patterns are possible (Dow, 2000). In the more interdisciplinary parlance of systems theory, economic systems are usually expressed as 'complex adaptive systems' which, while displaying complex and evolving dynamics, are often governed by mechanisms found repeatedly across different systems and disciplines (Holland, 1992). For these systems, then, it is ideal to have methods that are well-suited

to deal with the complexity of their openness, and simulations are one of the best tools for this purpose. Lawson (2009), based on his ‘critical realism’, actually advises that the complexity of social reality makes it entirely impervious to mathematical methods, arguing that any attempt is futile. However, most economists who subscribe to realism in economics, from the more theoretical Post Keynesians to the computationally-minded complexity and evolutionary economists, take a middle ground, arguing for ‘methodological pluralism’ (Colander, 2000). This means that a multiplicity of methods, including the standard modelling techniques of economics but also newer methods like simulations, ought to be used in a complementary way when conducting economic research of complex open systems.

Beyond general philosophical considerations, the specific characteristics of complex systems reveal how simulations can be useful. For instance, *indeterminacy* is a core facet of complex systems (Arthur, 2013). This refers to the inability to calculate the state of the system now and in the future, given the many variables and forces at play, and the inherent fundamental uncertainty that is unquantifiable and therefore impossible to calculate analytically. Part of the complexity of this system arises from a high degree of *nonlinearity* in the system’s dynamics, which could even be dangerous to model with traditional linear methods for the purpose of obtaining practical policy implications (Brock, 2000b). Nonlinearity and indeterminacy result in an *emergence* of patterns which come from the unique dynamics of the system in question, and which cannot necessarily be predicted via the observation of the system’s individual agents’ behaviour, as in micro-founded economic models (Harper and Lewis, 2012). This in turn creates an evolving or *adaptive* environment that can lead to *self-organisation* of agents, just like the emergence of like-minded cooperating factions in the games discussed in Section 2.4 (Holland and Miller, 1991). This all creates a system where, unlike traditional economic theory, history matters for understanding progress in the system, and processes are *path-dependent* (Foster, 2006). In sum, the characteristics of a complex economic system are vastly different to those of systems typically analysed in economic research, and accepting the relevance of complex systems in economics thereby necessitates consideration of methods like simulations that are better suited to deal with such complexity.


Another rationale for pursuing simulations is that the field is now supported by a robust range of methodological frameworks for simulation design and validation. Richiardi *et al.* (2006) noted that an earlier lack of ‘best practice’ guidance was one issue holding the field back, but as mentioned in Section 2.4, this began to change in the mid-2000s. Multiple articles and books arose, like Richiardi *et al.*’s (2006) common protocol for agent-based research, covering model structure, estimation/calibration guidelines, and validation checks; Szabó and Fáth’s (2007) detailed methodological guide on how to build simulations of economic games; Helbing and Balmelli’s (2011) economics-oriented tutorial on how to do agent-based simulations to capture emergent phenomena; Abdou *et al.*’s (2012) review of how to produce agent-based simulation research in a consistent manner, from identifying the research question to implementing the model and validating it; and Sargent’s (2013) interdisciplinary guide on how to verify and validate simulation models. Recently, Tesfatsion (2017) has consolidated much of this work into seven modelling standards for economists, ranging from agent definition and scope to system construction and the role of the modeller, as well as four empirical validation standards, covering input, process, descriptive output, and predictive output. Hence, by now there appears to be considerably more consistency in the approach to agent-based simulations than just 10 years ago, as the field has matured and general methodological inquiries have laid foundations for more advanced models (compare earlier *vs.* later simulations in Section 2.4), such that the computational economics field now seems well-justified, especially when

considering how the 2021 Nobel Laureate Symposium on Game Theory noted the need to explore numerical methods in addition to more traditional analytical methods. Simulations have been used by engineers, physicists, and other professionals for many years before they test their designs in the real world – now economists can follow suit.

It is important to recognise the limitations of simulations, but sometimes, perceived limitations can actually offer distinct advantages. Some economists rightly point out that simulations do not provide derived proofs of economic phenomena; instead, they offer examples, patterns, and exploratory findings that offer insights, but not certainties (Judd, 2005). Yet this precisely marks the distinction between *deductive reasoning* and *inductive reasoning*, which is vital to understand. Logical proofs offer a deductive certainty of their findings, but they require the aforementioned closure of a system to ensure no unaccounted forces can interfere. Inductive reasoning, on the other hand, seeks to make sense of patterns and trends, and admits their fallibility and possibility of change over time. When dealing with complex, open systems, deductive theorems are not always possible, and as the subject matter of this thesis deals exactly with a complex, dynamic economic system, simulations are perhaps the best tool to address related research questions. For example, consider the El Farol Bar Problem (Arthur, 1994): the El Farol bar in Santa Fe has a maximum capacity of 100 people, hosts a live music event every Thursday night, and the event is enjoyable if no more than 60 people attend. There is no rational, deductively solvable solution to win this game: if rational agents all expect no one will go, they will all decide to go, thereby overcrowding the bar and ruining the night; if they expect it will be overcrowded, no one will go, and it would have been fun after all. However, Figure 9 shows how simulations can inductively offer insights into how individuals actually decide over repeated Thursdays. The heterogeneous agents, who are endowed with differing amounts of heuristic rules to forecast attendance like ‘attendance will be the same as two weeks ago’ or ‘attendance will be the average of the last four weeks’ and develop preferences based on past performance, quickly self-organise into an ‘emerging ecology’. This results in them forecasting more than 60 attendees 40% of the time, and under 60 attendees 60% of the time, a pattern that is common in evolutionary ecology and economics. It is especially noteworthy that He *et al.*’s (2020) results from simulating the spatial public goods game with different parameterisations of persistence in reputation-based strategies in Figure 8 look remarkably similar to the results of the El Farol Bar Problem, with a clear equilibrium emerging around a certain level that cannot be calculated analytically, and that continues to fluctuate with some noise. Inductive reasoning methods like simulations are therefore useful when deductive analytical methods are ill-suited to solve a particular economic issue, or when the abstract nature of traditional economic methods are unable to model a sufficient degree of the issue’s real complexity.

By way of more general illustration, consider what is widely accepted to be a failing of traditional economic methods: the 2008 Global Financial Crisis (Buchanan, 2009). There were many causes to this crisis, and of course it is easier to declare in hindsight that it could have been predicted or prevented, but many of its failings arose from a poor understanding of complex adaptive systems. Risk models failed to see the illegitimacy of supposed risk diversification in financial derivatives markets; complex dynamics interlinking disparate parts of the economic network were missed or not adequately considered, leading to systemic risk and financial contagion; and emergent patterns including market bubbles and mass hysteria, driven by hard-to-predict network effects like herd mentality, arose with catastrophic results (Helbing, 2009). Since the financial crisis, agent-based methods have been specifically recommended to study large-scale macroeconomic systems over other approaches like DSGE models (Colander *et al.*, 2008), and central bankers are increasingly

Figure 9: *The El Farol Bar Problem: Simulating 100 consecutive Thursdays, it is clear that after a volatile start, agents inductively learn as a group to average around the 60-person attendance limit for an enjoyable evening out. Agents achieve this through heuristic rules of thumb, as opposed to perfectly rational game-theoretic decision-making. Reproduced from Arthur (1994, p.410).*



demonstrating an interest in agent-based modelling as an additional tool for understanding the complex macroeconomic and financial systems they are tasked with managing.

An important clarification is necessary: this thesis does not take the position that agent-based simulations and inductive reasoning are better than other methods in economics, be they econometric models or deductive analytical techniques. Rather, agent-based simulations are seen as an *additional* tool for exploring economic phenomena in *complementary* ways. As already discussed, simulations do not offer the deductive, logical completeness of mathematical proofs; researchers' licence to set initial conditions creates a need for significant robustness tests (Judd, 2005); and a good simulation study requires a strong rationale based on more theoretical or experimental research, such as to match a theory's results or to explore experimental findings in a more complex environment (Abdou *et al.*, 2012). Yet they do enable researchers to identify emergent patterns that cannot be deductively identified or proven; to explore dynamic behaviours in complex systems; and to model richly defined economic agents with detailed rules-based strategy updating mechanisms, going beyond simpler assumptions like perfect rationality. As such, when used in tandem with other available methods, simulations can offer additional support to the panoply of evidence for a more holistic economic theory or research field.

One open challenge with pursuing this method is that in the context of economics, it is not widely applied yet, and the simulations used to study behavioural economics theories are still developing in their capacities. However, as shown in Section 2.4, the literature has grown tremendously over the past 10 years and it has now reached a point at which interesting and nuanced insights are possible into how individuals cooperate in complex economic systems, beyond what can be assessed in experiments or with analytical methods like game theory. This thesis intends to further contribute to the field by more concretely uniting behavioural economics theories with the agent-based simulation method, and by exploring the robustness of previous findings in a more complex scenario of dynamically evolving public goods multipliers, as opposed to the simple fixed and positive multiplier of previous simulations.


3.2 Conceptual Framework

Ultimately, the objective of many agent-based simulations is to test theoretical results in an environment more closely approximating the real world (Tesfatsion, 2006), and this thesis aims to achieve exactly this. A wealth of behavioural economics research has demonstrated how cooperation exists to a greater extent than traditional game theory would suggest, and has proposed various mechanisms for controlling and encouraging cooperation, like punishment and reward systems. The model presented here will test some of these mechanisms in a more complex setting, namely, one in which many heterogeneous agents interact over a large number of rounds, and this section explains at a conceptual level how the model works.

3.2.1 Game Setup and Purpose

This thesis considers the strategic decision-making of a large number of economic agents in the canonical public goods game, to better understand how agents choose whether to cooperate in this social dilemma under more complex environmental features and dynamics than the traditional behavioural economics literature on public goods games. As this is an agent-based simulation, the agents are not real human beings like in experiments, but rather are modelled as cells on a spatial grid or lattice. In this form, the game is driven partially by the interactions between agents in close proximity to each other on the spatial grid, and the game is called the Spatial Public Goods Game (SPGG).

Figure 10: *The game's 2-D spatial lattice wraps around into a 3-D torus, in order to make sure all agents have 8 neighbours in the simulation.*



This setup is particularly useful for investigating how agents cooperate in a more realistic setting than previous experiments or other studies. Moreover, it permits the interaction of many agents – in this case, 2,500 agents distributed evenly across as 50x50 square lattice – for many rounds – in this case, 1,000 rounds – which is obviously not feasible in laboratory experiments due to time, money, and complexity issues. Moreover, the first 15 rounds of the simulation are calibrated to closely match experimental data from the recent public goods game literature, thereby setting a system of heterogeneous agents that nevertheless may act roughly as would be expected in the real

world. Also, the boundaries of the square lattice wrap around like a torus, so players on one edge of the square are actually neighbours with those on the opposite edge, thereby ensuring that every agent on the lattice has an equal number of cooperators (illustrated in Figure 10). As such, the SPGG allows for an exploration of cooperative dynamics in a complex adaptive system more closely mirroring the reality in which economic policies for promoting cooperation are applied to. This in turn can enable insights into unintended consequences of policies that might not be obvious from analytical evaluation alone, but might nevertheless emerge from the evolving complexity of the system; it provides a flexible sandbox environment for tweaking policies and testing how various mechanisms to encourage cooperation, like punishments or rewards, might fare in the context of these heterogeneous, realistic agents; and it offers an opportunity to test the robustness of findings on cooperation from behavioural laboratory experiments and other research in a more complex, realistic setting.

3.2.2 Agent Definition – Heterogeneity and Utility Functions

Once the behavioural game and simulation method have been determined, participating agents must be specified. To maintain close proximity to economic theory, agents in the game are defined as broadly self-interested, meaning that they update their strategies based on maximising their own utility, but their utility is also driven by other-regarding preferences to varying degrees. Critically for the study of complexity and realism of the simulation, agents are heterogeneous across multiple parameters such as utility function specification. This is important because significant heterogeneity can induce complex, chaotic dynamics, and Dreber *et al.* (2014) argue that heterogeneous agents could aid cooperation if they enable cooperation from those who were previously unsure or on the decision boundary between cooperating and defecting. Agents are also boundedly rational, and are imbued with myopic, adaptive learning rules to update their strategies based on the performance of other agents in their immediate neighbourhood, rather than using perfect information or foresight.

Defining agents' utility functions as being partly driven by other-regarding preferences was a logical outcome of designing a realistic public goods game. Two approaches were initially apparent: stick with the assumption of utility being driven by totally self-interested payoff maximisation akin to *Homo Oeconomicus*, or include other-regarding preferences in some way. In Section 2.2, it was highlighted that the perfectly rational *Homo Oeconomicus* agent is not very realistic, as it would imply that every agent defects in the game. Indeed, it does not withstand experimental evidence from behavioural economics, such as Fischbacher *et al.*'s (2001) finding that only around 30% of agents will always defect, or Zelmer's (2003) meta-study showing that in repeated public goods games, there is a consistent level of between 30-50% cooperation across rounds. Thus, other-regarding preferences must be included in agents' utility somehow.

Behavioural economists have generated alternative utility functions that incorporate some other-regarding preferences in order to reflect agents' desire to cooperate, such as *inequity aversion* (see Equation 2) or *conditional welfare* (see Equation 3). However, while these two utility functions are regarded as somewhat canonical in the behavioural economics literature, their effectiveness in reality has been brought into question, such as concerns around whether Fehr and Schmidt's cited data actually does support their theory (Binmore and Shaked, 2010). Attempts were made in this thesis to integrate these famous utility functions into the simulation, to try to capture a relevant degree of other-regarding preferences, but it was not possible to escape the zero-cooperation

trap seen with *Homo Oeconomicus* without using parameters drastically beyond the range typically used in the literature – effectively confirming the concerns of Binmore and Shaked about shaky empirical justifications. Furthermore, it would have been necessary to violate some of the parameter constraints stated in the original models (such as keeping $\alpha > \beta$ in the Fehr-Schmidt inequity aversion approach to ensure agents’ inequity aversion from opponents’ inequity does not exceed that from their own inequity), so these utility functions were not deemed appropriate for achieving a reasonable level of cooperation in the baseline simulation. Results for these two simulations are provided in Appendix A, but suffice it to say here that these models were not suitable for the spatial public goods game.

Hence, this model uses a more recently developed utility function offered by Alger and Weibull (2013) for *Homo Moral**is*, an agent whose utility is derived from a linear combination of both their own monetary payoff but also, in a nod to the famous Kantian ethical principle of acting only in the ways one would like to see everyone act, the payoff that would be achieved if every other agent also pursued their strategy. This utility function is expressed in Equation 7:

$$U_{\kappa}(x, y) = (1 - \kappa) \cdot \pi(x, y) + \kappa \cdot \pi(x, x) \quad (7)$$

for some degree of morality $\kappa \in [0, 1]$, where $\pi(x, y)$ is the purely monetary payoff that rational agents seek to maximise for themselves, and $\pi(x, x)$ is the morally-weighted payoff that would be achieved if all other agents ($j \neq i$) adopted the same strategy as agent i . Thus, when $\kappa = 0$, the agent is an entirely self-interested *Homo Oeconomicus*, and when $\kappa = 1$, the agent is an entirely virtuous ‘*Homo Kantiensis*’ (of whom Kant himself would be proud). The rationale for choosing this utility function is that it allows for cooperative preferences that can be easily calibrated to match experimental findings, and moreover, it has the best ‘hit rate’ for matching experimental evidence when compared to a range of possible utility functions (see Table 1 for comparative metrics). The final deciding factor for using *Homo Moral**is* is that it has only one free parameter, namely κ , whereas alternatives have two or more, so it is the easiest and simplest to calibrate without having to make judgments on which parameters to prioritise or how to manage trade-offs between them during calibration.

Table 1: Weibull et al. (2020) find that, out of six prominent utility functions, *Homo Moral**is*, *Reciprocity*, and *Conditional Welfare* specifications vastly outperform other specifications based on experimental evidence.

Utility Function	Hit Rate
<i>Homo Oeconomicus</i>	28%
Unconditional Altruism	40%
Inequity Aversion (Fehr and Schmidt, 1999)	60%
Conditional Welfare (Charness and Rabin, 2002)	82%
Reciprocity (Charness and Rabin, 2002)	83%
<i>Homo Moral</i> <i>is</i> (Alger and Weibull, 2013)	83%

3.2.3 Agent Definition – Strategies and Updating Rules

With agents' utility functions defined, their strategies and strategy updating rules can next be constructed. Given the nature of the public goods game, agents have only two strategies in the base case: cooperate, thereby paying into the public goods pot, or defect, which means not paying into the public goods pot. To differentiate between these two strategies, cells on the spatial lattice occupied by cooperators are coloured blue, while cells with defectors are coloured red. In later iterations of the SPGG with social mechanisms like punishment and reward, these additional options are represented by further colour-coding (e.g. yellow cells for punishing cooperators, green cells for rewarding cooperators, etc.). Hence, the strategy of every agent can be shown every round just by looking at the spatial lattice, and observing the spatial lattice over multiple rounds can reveal emergent patterns and spatial dynamics such as regularly evolving strategies and more resilient clusters of the same strategies.

Evolutionary dynamics are of great importance in this agent-based simulation. In fact, evolutionary game theory is the foundation for agent-based simulations of behavioural games (Szabó and Fáth, 2007), and initially entered economics via evolutionary biology as a way to identify Nash equilibria arising from evolutionarily dominant game strategies (Banerjee and Weibull, 1996). As noted in Sections 2.3 and 2.4, evolutionary biologists have been influential in guiding the SPGG literature, and to this day the most common strategy updating rules for agents in the SPGG are *imitation* and *innovation*, both of which are grounded in evolutionary biology. So as to remain methodologically close to the literature, this thesis therefore includes these two learning rules as the updating rules for agents to re-evaluate their choice of strategy after each round of the game.

Innovation is a strategy updating rule more overtly borrowed from the evolutionary biology literature. Much of the field of evolutionary biology is based on the idea of occasional random mutation, which introduces new behaviours or survival strategies into a species. This concept of mutation has been adapted in the context of SPGGs to represent a small probability that agents randomly switch their strategy to an alternative, such as randomly switching from cooperation to defection for no rational, utility-driven reason. If the strategy is a good option for the agent, it will likely remain that agent's strategy for a while. However, due to the uncertain, evolving nature of a complex adaptive system like this, the point is not necessarily that a successful mutation will become a long-term evolutionarily stable strategy, as in some of the more traditional evolutionary game theory research, because it could ultimately be displaced by a changing environment or other dynamically evolving parameters. On the contrary, the strategy may only remain successful for a short while. Yet this sort of 'irrational' changing of strategy is often detected in reality: it could be that agents make a mistake, or that they simply want to change their strategy to see what happens (Ariely, 2010). It is thus a small but realistic and powerful strategy updating rule, that introduces a potentially large element of complexity in the system's dynamics and adds an important concept in the study of evolutionary processes more generally.

Imitation is the more common strategy updating rule that agents use, and it can be specified in a few different ways. The version used here states that after every round, each agent has a probability of imitating the strategy of their best performing neighbour. Their neighbours are the eight agents directly neighbouring their cell on the spatial lattice, which is known as a Moore neighbourhood. This means that agents learn from their immediate surroundings, and as their surroundings evolve, their strategies adapt accordingly. The advantage of using probabilities instead of a strict rule of

‘imitate-if-better’ is that it removes the property of payoff monotonicity, meaning that agents do not always choose strategies that seem to make them better off, which is an important component of bounded rationality (Szabó and Fáth, 2007) and is more conducive to realism. In this model, the agents are myopic and thus only compare the performance of neighbours after the most recent round, although this could be extended to include a memory of a few rounds of performance and still be reasonable from the perspective of economics (Bordalo *et al.*, 2017). Further, since the strategies of the next period are completely determined by data from the last, the system becomes a Markov process. Hence, the agents are in aggregate essentially estimating a Markov chain, and the entire process can then be implemented via Markov Chain Monte Carlo simulation.

The probability of imitating an agent’s best performing neighbour is determined by the Fermi-Dirac equation, which is shown in Equation 8. The equation compares the utilities of the agent (U_i) and their neighbour ($U_{j \neq i}$) in a way such that if they have equal utility, the probability of imitating their neighbour (that is, strategy $s_i \rightarrow$ strategy s_j) is 50%. However, as the neighbour’s utility tends towards being infinitely larger (smaller) than the agent’s, the probability of imitating tends towards 100% (0%). Further, a noise parameter K is included, which distorts the imitation probability to represent the fact that it is not always clear to an agent whether their neighbour is better off than them. A greater K reduces the probability of an agent switching strategy even though it would result in a utility gain, and vice versa. The origin of the equation lies in statistical physics, but it is commonly used in the SPGG literature for this purpose (e.g. Zhu *et al.*, 2020a).

$$IM(s_i \rightarrow s_j) = \frac{1}{1 + e^{(U_i - U_j)/K}} \quad (8)$$

3.2.4 Game Mechanisms

The last feature of the models in this thesis are the additional social mechanisms that are introduced in an attempt to improve the proportion of cooperation. At a conceptual level, these include:

- *Punishment*: While agents continue to play the public goods game at the global level – that is, cooperators pay into a global public goods pot that is distributed across all 2,500 agents – they may be dissatisfied with agents choosing to defect within their immediate vicinity of 8 neighbours, and can therefore choose the additional strategy of punishing those neighbours while still cooperating. Thus, if an agent opts for this strategy, they incur a cost to themselves in order to punish the defectors directly adjacent to the punisher. The cost is proportional to the amount of surrounding defectors, and the punishment received by the defector is usually higher than the cost to punish them. This differential between cost of punishing and fine levied is typical in the punishment literature, but a simple example of a punishment being higher than the cost of punishing is destroying something that the defector values – the act of destruction does not cost the punisher much, but could cause a great loss of value to the targeted defector. However, if there are no neighbouring defectors then the punisher suffers no cost and no punishment is dealt by that agent. This introduces an interesting dynamic whereby cooperators benefit from the presence of punishers, but may choose to free-ride on punishment as well (called second-order free-riding in the literature).
- *Reward*: Equally, cooperators can try to incentivise more cooperation in the game by choosing

the strategy of rewarding neighbouring cooperators. Again, the reward mechanism is set up with the cost of rewarding typically being lower than the reward granted to cooperators, as in the case of punishment. Praise, which is inexpensive, and small gifts are two examples of rewards that can provide a higher utility benefit to the recipient than cost incurred to the gift-giver. The mechanism works just as punishment does, since rewarding is local and thus the cost of rewarding is dependent on whether or not the agent has any neighbouring cooperators, and the cost is proportional to the amount of cooperators surrounding the agent.

- *Policing*: In contrast to the punishment mechanism, which induces punishment from local dissatisfaction with defectors, a policing mechanism operates on a global level, which is to say that cooperators who want to punish defectors can pay to support a global police force, which then levies punishments across all defectors on the spatial lattice. Again, it is reasonable to assume that the policing fine levied on defectors can be higher than the cost of policing, since fines can be set almost arbitrarily high by the judicial system, and high fines are generally known to deter unwanted behavior. This provides a similar mechanism to punishment, but varies the method of implementation from local to global and is designed to provide additional insights into how such factors can influence the effectiveness and robustness of punishment mechanisms more generally. One caveat is that policing is a less common mechanism in the literature, though Liu *et al.* (2018) adapt a synergistic punishment mechanism to mirror policing as mentioned in Section 2.4, and Dong *et al.* (2016) make a similar distinction as this thesis does between individual incentive systems like punishment and reward on one hand, and institutional incentives like policing on the other hand.
- *Reputation*: The reputation mechanism is implemented to make agents care about how they are perceived by others, and this image is impacted solely by whether they are cooperating or not. It is implemented in a similar way to Zhu *et al.* (2020a), by changing the agents' utility function. This is done by introducing a reputation parameter R as a scalar for the old utility function, such that any utility the agent would have previously is now scaled by their reputation, as can be seen in Equation 9:

$$U_{i,t}(x, y) = R_{i,t} * ((1 - \kappa) \cdot \pi(x, y) + \kappa \cdot \pi(x, x)) \quad (9)$$

The reputation scalar R increases (decreases) when agents cooperate (defect) by a constant value of θ , as seen in Equation 10 below. Also, R can take values in the interval of 0 and 2. Thus, when an agents' reputation is zero they receive no utility and agents with maximum reputation receive double utility.

$$R_t = \begin{cases} R_{t-1} + \theta & \text{if cooperated,} \\ R_{t-1} - \theta & \text{if defected,} \\ 1 & \text{if } t = 0. \end{cases} \quad (10)$$

- *Moral Adjustment*: The final mechanism considered is called 'moral adjustment', and does not feature noticeably in the literature to date. It is instead designed specifically to assess how one of the novel features of this SPGG model, the *Homo Moralis* utility function and its κ parameter, could be addressed to encourage more cooperation. Effectively, the moral adjustment mechanism works in a similar way to the reputation mechanism, but rather than directly influencing the monetary payoff of agents, it instead influences their unique κ

parameter value. When agents cooperate in a round of the game, this is deemed to imply a preference to choose the ‘right’ moral strategy, and upon success from cooperation this influences the underlying moral preferences of the agent, increasing their κ parameter positively by a pre-specified amount $\Delta\kappa$. Conversely, if a choice to defect results in relative success, this teaches the agent to reduce their κ parameter by $-\Delta\kappa$, reflecting a shift towards more selfish or less ‘moral’ preferences. Throughout, the κ parameter remains bounded between $0 \leq \kappa \leq 1$, as per the original theory.

3.2.5 Game Procedure and Outcomes

With all components of the game now specified, attention can be turned to the baseline simulation’s game procedure. At the beginning, agents are randomly distributed across the 50x50 square lattice with equal likelihoods of cooperating or defecting. As randomisation is common in laboratory experiments, this is a reasonable starting point for the simulation too. Then, they play the first round of the SPGG, with cooperators chipping into the public goods pot, the contents of the pot being multiplied by the multiplier, and this magnified pot being redistributed across all agents. Once payoffs have been allocated, the utility of each agent can be calculated based on their choice of strategy, their payoff, and their unique *Homo Moralis* utility function parameter, κ . For example, even if a cooperating agent receives a small payoff due to a large number of defectors, they may still obtain a higher utility if their κ parameter is large and they mostly receive utility just from the moral component of their act of cooperation.

After utilities for all agents are computed, it is time for strategy updates. First of all, each agent undergoes a small probability of their strategy mutating randomly, as per the ‘innovation’ strategy updating rule. For the vast majority who do not innovate, the agents then look around their eight closest neighbours, *a.k.a.* their ‘Moore neighbourhood’, and identify the highest-utility agent around them. The main alternative to this Moore neighbourhood is the diamond-shaped von Neumann neighbourhood, with four neighbours at North, East, South, and West of a particular agent; while a much larger neighbourhood size of *e.g.* 24 can increase clustering dynamics (Zhu *et al.*, 2014), Szabó and Fáth (2007) note that both the Moore and von Neumann neighbourhoods are common options in the SPGG literature, and in this case the more information-dense Moore neighbourhood is chosen. If the highest-utility neighbour has a higher utility than the agent’s own, they undergo a probabilistic chance of updating their strategy to match that of their highest-utility neighbour, where their probability of updating is based on the difference between their utility and their neighbour’s (as per the Fermi equation). With a new set of chosen strategies for the next round, each agent updates the colour of their cell on the spatial grid if they have changed strategy, and the next round of the game begins. The game continues to iterate round by round via the same procedure, and after the final round the simulation produces visualisations of results, in particular the proportion of cooperation over time, snapshots of the spatial lattice at various time intervals during the procedure, and time series of strategy payoffs and average agent utility throughout the game.

Once this baseline simulation has been run, the entire procedure can be followed again and again in the Monte Carlo fashion with the addition of social mechanisms for encouraging cooperation. Some of these mechanisms may include entirely new strategies, such as introducing the distinction between punishing cooperators and non-punishing cooperators, while other mechanisms may act only indirectly on the level of cooperation by influencing agents’ preferences over the two baseline

strategies, cooperate and defect. Finally – partly to test robustness of different mechanisms, and partly to increase the complexity of the system further – simulations are re-run with one extra change: the SPGG multiplier is no longer a fixed value, but rather fluctuates in semi-stochastic, semi-cyclical manners that more closely resemble the uncertain decision-making environment of many real-world economic situations. Given agents' myopic, adaptive learning rules, this introduces another layer of unsolvable, unpredictable mechanics that present an even more chaotic environment for determining the dominance of different strategies, and potentially will proffer further insights into the nature of cooperation in social dilemmas and the effectiveness of different social mechanisms.

4 Model

This section briefly explains how the simulation model was implemented in MATLAB, before explaining in detail what relevant data sources and calibration methods were used to specify the model. Finally, pseudo-code outlines the game procedure as it is followed in the simulation code.

4.1 MATLAB Implementation

In practice, this game can be implemented using many different programming languages, and there is no one language that dominates the SPGG literature. All simulations conducted in this thesis were constructed from scratch via MATLAB scripts and functions. MATLAB was the programming language of choice due to its high degree of flexibility in creating sandbox simulations, its applicability in running Monte Carlo simulations using large arrays to capture variables, and its significant popularity and familiarity within other fields of quantitative economics.

The main script file used contains the baseline simulation and a number of flexible design parameters such as whether agents possess strict *Homo Oeconomicus* utility specifications or some heterogeneous distribution of *Homo Moralis* utility specifications. A few key processes are written as separate function files to be called upon in the main simulation code, such as the functions for identifying the highest-utility neighbour for each agent in each round, but otherwise options are selected within obviously flagged segments of the main code file itself. This is because after the baseline simulation was produced, five separate simulation code files were created for each of the five social mechanisms being tested, so splitting every process in the simulation into a separate function file would have led to a confusingly high number of associated files. Further details as to specific scripts and functions can be found in Appendix B.

4.2 Data

Agent-based simulations of spatial public goods games are not particularly data-intensive, but simulations do require rigorous justification of choices when calibrating parameters and initial conditions. In fact, due to the path-dependent nature of complex systems more generally, diligent care must be taken in selecting these initial conditions, or else the system might collapse into an unrepresentative mess of dynamics. Thus, two categories of data are used to program this simulation.

4.2.1 Identifying Parameter Values from Previous Research

The first category of data used is prior research papers on spatial public goods games and other economic concepts within the model. These provide many insights into relevant parameter choices for the simulation, and it is standard practice in computational economics to use tried and tested parameters from prior studies. Moreover, most parameter values are borrowed from prior agent-based simulations or related work from recognised academics in the field, and are derived from reputable journals including *Games and Economic Behavior* and *Science*. κ is the one remaining parameter after these borrowed parameters, and is calibrated in Section 4.2.2.

Table 2: Game parameter values in the baseline simulation model and their respective sources.

Parameter	Value	Source
Chance of ‘Innovation’	0.1% per round	Hauert <i>et al.</i> (2007)
K (in Fermi Equation)	0.1	Amaral and Javarone (2018); Zhu <i>et al.</i> (2020a)
κ	$\sim U(0.1, 1.0)$	Calibrated (see Section 4.2.2)
Multiplier	1.6x	Lugovskyy <i>et al.</i> (2017)

Table 2 outlines all of the set parameters in the baseline simulation. Starting at the top:

- *Chance of ‘Innovation’*: The chance of ‘innovation’ in strategy, *i.e.* random switching (or ‘mutation’ in evolutionary biology) is set to 0.1% per agent per round, in line with that set by Hauert *et al.* (2007), a group of agent-based simulation experts including the aforementioned Traulsen, Nowak, and Sigmund in a study on costly punishment in the public goods game.
- *K (in Fermi Equation)*: This parameter regulates the degree of noise or uncertainty in the Fermi Equation, which governs agents’ strategy updating role based on imitating more successful neighbours. This is the standard method of imitation learning in SPGGs, and as mentioned in Zhu *et al.* (2020a), $K = 0.1$ is the standard parameterisation across the literature.
- κ : This parameter measures the degree of ‘morality’ in Alger and Weibull’s (2013) *Homo Moralis* utility function, and takes a uniform distribution bounded between 0.1 and 1.0. Section 4.2.2 explains how this parameter is calibrated to closely resemble representative experimental data.
- *Multiplier*: The multiplier, also referred to in net terms in the literature as the Marginal Per Capita Return (MPCR), is the amount by which the public goods pot is multiplied each round. Metastudies reveal that this MPCR is usually in the range of 0.3-0.6, meaning that the gross multiplier in this game ought to be within 1.3x-1.6x for adequate comparability with typical experiments (Cartwright and Lovett, 2014). In this thesis, a multiplier of 1.6x is used as it matches that used in the experiment to which κ is calibrated (see Section 4.2.2).

In summary, all parameters for the baseline simulation except for κ are taken from reputable and relevant sources in the literature. The next section explains how the final parameter, κ , is calibrated such that the baseline simulation is validated by relevant experimental data. In addition to the baseline simulation’s parameters, social mechanisms introduce the additional parameters in Table 3. These parameters take on a range of possible values, as multiple specifications of the mechanisms are tested to compare them. These are justified as follows:

- *Punishing Cost and Fine*: To get a better understanding of how the impact of the mechanism changes with different specifications, a range of both high and low punishment fines and costs are considered. Cases where the punishment and cost were equal are also considered, as well as where the fine is double the cost. Zhu *et al.* (2020b) use a fine of 0.5 and a cost of 0.05, while Helbing *et al.* (2010) study costs and fines between 0-100% of per-round endowment,

Table 3: Parameter values for additional social mechanisms.

Parameter	Value Range
Punishing Cost	0.05-0.1
Punishment Fine	0.1-0.4
Rewarding Cost	0.05-0.1
Reward Value	0.1-0.4
Policing Cost	0.0625-3.125
Policing Fine	0.1-5
Reputation Δ	$\pm 0.1-0.2$
Reputation Multiplier Minimum	0-0.5
Reputation Multiplier Maximum	1.5-2
Moral Adjustment Δ	$\pm 0.001-0.05$

and argue that punishment that is too costly compared to the fines will make punishment a non-viable strategy, further motivating the exploration of multiple specifications.

- *Rewarding Cost and Value:* As punishing and rewarding are two antithetical yet similarly structured strategies, this thesis opted to use equal values of rewards and their costs for comparative purposes. In the literature, Szolnoki & Perc (2010) use rewards between 0.1-1.5 and costs between 0.01-0.2, placing the rewards in this thesis in the lower range of their values and costs in the middle of their range.
- *Policing Cost and Fine:* As the policing mechanism is more exploratory with less available literature, the policing fine is set to vary from a very small level of 0.1 to a medium-high level of 5, to better enable exploration of how different fines encourage or discourage cooperation. In this setting, the cost of policing has been set as $Cost = \frac{Fine}{Multiplier}$, to mirror how a institutionally driven, globally (or nationally) provided service like policing is likely to benefit from leverage effects in the same way that government services often benefit from a government multiplier effect.
- *Reputation Multiplier and Δ :* The main paper which discusses reputation in the spatial public goods game in a context similar to this thesis is Zhu *et al.* (2020a), which uses a reputation multiplier range between 0 and 2 and a reputation increment that varies between 0.1 and 1. This increment was deemed too extreme in this model, however, so instead only values that are at the lower end of this interval are used (specifically, 0.1-0.2). Another area of interest in this thesis is what happens with a less extreme range for the reputation parameter, so an alternate specification with a possible reputation interval between 0.5 and 1.5 is also used.
- *Moral Adjustment Δ :* This round-by-round increment for adjusting the morality parameter κ in the *Homo Moralis* utility function is designed to be very small, as low as 0.001, because κ can only vary between 0 and 1, so the incremental change has to be very small in order to not reach extreme results too quickly. This is the only mechanism that is entirely novel, having been created by the authors to allow for dynamism in the also recently-developed *Homo Moralis* utility function specification, so in this instance the authors selected the parameter values at their discretion, and it is emphasised that results from this social mechanism are less important than results from other mechanisms.

4.2.2 Calibrating Parameter Values from Previous Experiments

The second category of data used is experimental outcomes from repeated public goods games, which are primarily used as a way to validate the choice of the one remaining free parameter in mode, namely the *Homo Moralis* utility function parameter κ . Due to the early stage of research *Homo Moralis* as a utility function, there is not a significant or reliable literature available with estimated parameter values for κ , so this thesis calibrates κ in such a way that the first 15 rounds of the baseline simulation closely match levels of cooperation found in experimental games. If the simulation does indeed resemble experimental results, this offers strong validation for the parameters and initial conditions chosen.

Reviewing the earlier public goods experiment literature, prominent behavioural theorists Dawes and Thaler (1998) noted 40-60% cooperation in a one-round game, and in finitely repeated games contributions start around 53% but decline to below 20% within five rounds. Across finitely repeated games with alternative contexts, there appears to be broad acceptance that the typical proportion of cooperation lies somewhere between 30-50%: Zelmer’s (2003) meta-analysis of public goods games reveals an average contribution of 37.7%, while Fiala and Suetens’s (2017) more recent meta-study of cooperation with different levels of transparency in information finds an average of 40% in cooperation. On a round-by-round basis, Andreoni and Croson (2008) find that strangers in finitely repeated games start between 40-50% cooperation and decrease linearly as the game progresses, while partners start at around 50-60% cooperation and decrease to a lesser extent. Similarly, Kawamura and Tse (2021) find that when splitting game participants into high and low cognitive ability, both high and low cognitive ability participants start cooperating around 35-50% depending on the treatment and decline to 30% or lower, with lower final cooperation levels in the low cognitive ability groups. It therefore appears that, on average, contributions starting around 50% and declining to between 30-40% is the typical trend in repeated public goods games.

Most relevantly, though, is the investigation of round-by-round contributions in ‘infinitely’ repeated games, as these allow for more precise calibration of the simulation to experimental data. Lugovskyy *et al.* (2017) provide exactly this in the context of the public goods game, basing their methodology on similar, widely accepted prisoner’s dilemma experiments by Dal Bó (2005) and Dal Bó and Frechette (2011). In these experiments, behaviour under infinite repetition is induced by attaching a high probability that after each round there will be another round, and cooperation levels are reported for each round. In Lugovskyy *et al.* (2017) the per-round contributions with an MPCR of 0.6 are approximately those in Table 4, and building on this, the baseline simulation is calibrated with κ being uniformly distributed across a lower bound of 0.1 and an upper bound of 1. Comparisons between the experimental data and potential parameterisations are provided in Table 5.

In this case, both experimental data and the baseline simulation dip downwards from the 50% mark, and from there onward oscillate within a range of about 30% to 50%, ending together around 35%. The numbers do not match precisely for every round – this would likely be a sign of overfitting anyway – but there is a high degree of similarity in specific levels of cooperation, such as the fluctuation from around 48% in round two, to 30% in round three, and back up again. Experimental data is slightly more volatile and irregular than the simulation, but this is not of major concern as both the levels of cooperation and volatility of the series are consistently similar, and the Mean Absolute Error of their difference is minimised with $\kappa \sim U(0.1, 1)$. Finally, it is of great interest that this closely matches Weibull *et al.*’s (2020) estimation of a reasonable range for κ , given Weibull’s

role in creating *Homo Moralis*.

Table 4: Approximate proportions of cooperation (ρ_c) from the first 15 rounds of a representative public goods experiment with infinite-like repetition and a 1.6x multiplier (taken from Lugovskyy et al., 2017).

Round	ρ_c (Lugovskyy et al., 2017)
1	37%
2	47%
3	33%
4	46%
5	38%
6	40%
7	28%
8	21%
9	39%
10	48%
11	38%
12	45%
13	39%
14	52%
15	35%

Table 5: Comparing the proportion of cooperation (ρ_c) in the first 15 rounds of experimental data against the baseline model with three different calibrated uniform distributions of κ parameters, the range of $\kappa \sim U(0.20, 0.95)$ displays a close similarity in certain rounds, while the range of $\kappa \sim U(0.1, 1)$ is selected for the baseline model due to its lowest Mean Absolute Error.

ρ_c Round	Experiment	$\kappa \sim U(0, 1)$	$\kappa \sim U(0.1, 1)$	$\kappa \sim U(0.2, 0.95)$
1	37%	18%	27%	27%
2	47%	43%	52%	49%
3	33%	26%	29%	29%
4	46%	47%	53%	49%
5	38%	25%	28%	31%
6	40%	47%	51%	50%
7	28%	26%	30%	31%
8	21%	45%	51%	50%
9	39%	27%	30%	32%
10	48%	45%	51%	51%
11	38%	28%	31%	33%
12	45%	44%	51%	49%
13	39%	29%	31%	33%
14	52%	43%	51%	49%
15	35%	28%	32%	34%

4.3 Game Procedure

Following the game procedure outlined in Section 3.2.5, a more specific overview of the game procedure is provided here, in the form of pseudo-code:

SETUP

```
Initiate spatial lattice with randomised agents;  
Define initial parameters and state variables;  
Set Homo Moralis parameters and distribution across agents;  
Set multiplier type (fixed; dynamic – random walk,  
dynamic – cyclical);
```

GAME

```
Start game with For Loop (loops for every round);  
    Compute agents' pre-round contributions, punishments,  
        etc.;  
    Compute agents' post-round payoffs and utilities;  
    Update strategies driven by random innovation;  
    Update strategies driven by imitation of highest-  
        utility neighbour;  
    Update dynamically evolving parameters like dynamic  
        multipliers (where applicable);  
    Compute proportion of cooperation;  
End For Loop
```

PLOT CHARTS

```
Plot proportion of cooperation over time;  
Plot strategy payoffs over time;  
Plot utility per capita over time;  
Plot snapshots of lattice over time (at T = 1, 3, 5, 10,  
15, 20, 30, 50, 100, 200, 500, 1000);
```

END

It is this broad procedure that the MATLAB code implements, with dynamically evolving parameters and all results being stored in arrays to track each parameter for each agent for every round. Further insights into the scripts and functions are available in Appendix B, and the MATLAB code is available in a related supplement to this thesis.

5 Results

5.1 (Overly) Simple SPGG with *Homo Oeconomicus*

Briefly, this first model setup is presented to show why prosocial preferences must be specified for the agent-based simulation method if it is to obtain valid results mimicking experimental data. In this model, agents display the purely self-interested focus on monetary payoffs of *Homo Oeconomicus*, and play the public goods game 1,000 times. No social mechanisms, like punishment or reward, are included yet, to aid in comparison with experimental public goods games. In this SPGG, agents can therefore only be cooperators or defectors. They follow the typical game procedure of starting in randomised locations, playing the first round, and updating their strategies based on innovation (with 0.1% probability) or otherwise based on imitation, driven by the relative difference between their utility and that of the highest-utility neighbour in their immediate neighbourhood. As can be seen in Figures 11 and 12, the strictly higher payoff that arises from defection leads to a decisive shift towards defection for every agent, resulting in a long-term equilibrium of almost 100% defection from the very beginning. Total defection is only prevented by the occasional random mutation or innovation of strategy, and agents who do undergo this mutation quickly re-learn to defect. It is clear that with purely selfish preferences focused on monetary payoffs and no punishment or other mechanisms, the typical experimental outcome of 30-50% cooperation illustrated previously is not possible.

Figure 11: Evolution of the SPGG with only *Homo Oeconomicus* agents over 1000 rounds (blue agents are cooperators, while red agents are defectors). Note that due to the strength of pure self-interest and focus on monetary payoffs, the only resistance to the dominant defection strategy is a very meagre pushback from random innovation of strategy.





Figure 12: With only homogeneous *Homo Oeconomicus* agents, the strictly dominant payoff from defecting rapidly gives rise to approximately 0% cooperation for the remainder of the game.




5.2 Baseline SPGG with *Homo Moralis*

It was made clear in Sections 2.2-2.4 and 5.1 that *Homo Oeconomicus* is not a sufficient characterisation of cooperative decision-making to match the real world. Other-regarding preferences are an important driver of such decisions, and the Kantian moral imperative to do as one would hope others in society to do is arguably a much more realistic motivation for cooperating. By way of illustration, consider the macro-scale need for cooperation against climate change, which requires a significant majority of people to reduce their personal emissions in their daily lives, despite there being no clear monetary incentives to do so. It is obvious from general observation that there are many individual agents who do take emissions-reducing actions, so *Homo Oeconomicus* is not the ideal model for characterising this behaviour. These agents also know their actions are only meaningful if many others cooperate in a similar way, as per the Kantian moral imperative built into the *Homo Moralis* utility function. Yet there are still plenty of people who free-ride on the emissions-reducing behaviour of others in a way similar to *Homo Oeconomicus*, so cooperation is by no means perfect. As such, being able to characterise economic agents' cooperative dynamics with the *Homo Moralis* utility function and a wide distribution of the morality parameter κ enables an economic researcher to capture both the self-interested *Homo Oeconomicus* and more morally conscious *Homo Moralis* agents in one expression. Having already calibrated this in Section 4.2.2, the following results reveal how this approach more closely approximates reality in the public goods game and sets a powerful baseline upon which social mechanisms can be assessed via simulation.

Observation 1: Having calibrated the model to fit experimental data for the first 15 rounds, the model continues to display expected stylised facts from the literature, namely that cooperation oscillates roughly around $\rho_c = 40\%$ and in the range of approximately 30-50%.

Figure 13 reveals game dynamics that appear remarkably consistent with experimental data from Lugovskyy *et al.* (2017) and stylised facts from meta-studies of public goods game experiments from Zelmer (2003). The average proportion of cooperation over the entire game is about 41%, and this

Figure 13: With a distribution of heterogeneous *Homo Moralis* agents, average cooperation reaches around 41% across all rounds. Initial rounds of the game closely resemble outcomes from real public goods game experiments, and the remaining rounds validate the stylised fact that cooperation fluctuates between 30-50%.



proportion oscillates back and forth roughly between the range of 30-50%. It is clear that other-regarding preferences are a modelling necessity to produce results that realistically approximate actual human behaviour, and the *Homo Moralis* utility function offers a parsimonious way to obtain such preferences. With only one free parameter – the morality parameter κ – and a reasonable assumption that agents are heterogeneously endowed with unique κ values, a richly complex game environment is created. Overall, this model appears to present a robust and realistic baseline, on top of which social mechanisms can be constructed and explored.

***Observation 2:** After the first 100 rounds, the range of ρ_c oscillations narrows, before widening again for most of the game. Then, roughly between rounds 700 to 950, the game enters an extremely stable ρ_c range tightly around 40%, before suddenly widening again just before the end of the game.*


Beyond model specification, Figure 13 offers plenty more of interest. While experimental data and stylised facts are met, the specific pattern of the proportion of cooperation, ρ_c , throughout the game is highly irregular and volatile. It is a fairly common pattern to see narrowing of oscillations around the average for the first 100 rounds in the agent-based simulation literature, as this is the time over which agents learn from their neighbours by imitating, settle into defensive clusters, and better define what strategies make sense given the behaviours of their surrounding opponents. Then, for much of the game, a fairly stable range of oscillating ρ_c is followed, representing a sort of sustainable equilibrium dynamic, which is interesting to see as it validates the finding in meta-studies that cooperation in repeated public goods games oscillates between 30-50%, while doing so over hundreds more rounds than the experimental literature can assess. Moreover, it suggests that there is a long-term equilibrium of approximately 40% cooperation (absent any social mechanisms to

improve cooperation), which is revealing about how cooperative engagements are likely to progress without influence – for example, if only 40% of people naturally take action to mitigate their contributions to climate change, it will not be possible to achieve climate change targets by taking a *laissez-faire* approach, and this in turn highlights the importance of social mechanisms to influence more cooperation via public policy, regulation, or other mechanisms.

Observation 3: *Spatial interactions are important drivers of cooperation across the game lattice, with defensive clusters becoming visible from an early stage of the game and holding similar patterns for most of the game after $T = 10$.*

Figure 14 is particularly useful in revealing how clusters of cooperators and defectors are formed and changed over time, as denoted by their respective colours of blue and red. Both the defensibility of these clusters and their speed of development are relevant: they form a clear, consistent spatial distribution of clusters within the first 10 rounds of the game, which is remarkably fast, although is likely driven somewhat by the relatively simplistic imitation-driven strategy updating rule these agents are imbued with. These clusters roughly hold their shapes throughout the entire game after $T = 10$, demonstrating how strongly the local social ties between neighbours can shape the

Figure 14: *Evolution of the SPGG with a distribution of heterogeneous Homo Moralis agents over 1000 rounds (blue agents are cooperators, while red agents are defectors).*




behaviour of all agents over the course of the game. An implication in the climate change context might therefore be that agents are quick to settle on their decision to cooperate or not for the long term, and social communities might be an effective medium for influencing defectors to cooperate – or conversely, in the case of defecting clusters, a tricky barrier to overcome – in reducing individual emissions for the social optimum.

Observation 4: *The Homo Moralis baseline simulation creates a higher utility per capita across the entirety of the SPGG, compared to the pure Homo Oeconomicus simulation, which sees an immediate and irreversible drop to approximately 0 for the whole game.)*

In Figure 15, the left-hand panel shows how the pure *Homo Oeconomicus* simulation effectively results in an immediate and permanent drop to 0 utility, which is a bad outcome for all, and is driven by the fact that with no cooperation, there is no possibility for anyone to earn a positive payout from the public goods multiplier. The right-hand panel shows considerably higher utility per capita in the baseline model with a heterogeneous distribution of *Homo Moralis* agents, and notably this utility per capita varies positively with the proportion of cooperation (compare this chart with that in Figure 13). This offers further evidence that encouraging cooperation is beneficial for society, and highlights how suboptimal the *Homo Oeconomicus* mentality is for dealing with real-world cooperative dynamics. These utility charts will not be repeated for every social mechanism, but it is worth noting that this positive relationship between cooperation and utility per capita remains consistent across all mechanisms.

Figure 15: *A comparison of total group utility over time from the simple Homo Oeconomicus SPGG simulation (left-hand side) and the baseline Homo Moralis SPGG simulation (right-hand side).*



5.3 SPGG with Punishment Mechanism

Observation 5: *With punishing cooperation as an added strategy, the proportion of cooperators no longer oscillates rapidly and the average level of cooperation can increase by 8-20 percentage points, depending on the strategy parameters. All specifications also show the pattern of an initial spike and dip in cooperation, followed by a positive trend until the end of the game.*

Figure 16: Proportion of cooperation (ρ_c) over time in the SPGG with punishing cooperation as an added strategy and with different punishment fines and costs.




Figure 16 shows that adding punishing cooperation as a strategy for agents has a moderating effect on the fluctuations in the level of cooperation in the game. In the baseline simulation (Figure 13), there was a great deal of noise around the average level of cooperation, compared with almost zero noise here. Varying the punishment fines and costs as in Table 6, the effect on average levels of cooperation is always positive and the increase relative to the baseline of roughly 40% average cooperation varies between 8-20%. This table also shows that a higher ratio between the fine and the cost increases average cooperation further, but in the cases of equal ratios, the scenario with a higher fine and cost favours cooperation even more. This could be because a higher absolute level of punishment causes fewer punishers to be needed to eliminate the monetary advantage of defecting. The trend of increasing cooperation might be explained by a growth in punishing cooperators, which could start a feedback effect that favours cooperation and deters defection.


Table 6: Average proportion of cooperation (ρ_c) by specification of punishment mechanism parameters.

Fine	Punishment Cost	Average ρ_c ($T = 1000$)
0.2	0.1	52.26%
0.4	0.1	60.49%
0.2	0.05	54.66%
0.1	0.05	48.25%

Observation 6: The clustering effect observed in the baseline scenario remains in a meaningful way. Also, the amount of punishing cooperators outnumbers the amount of regular cooperators in the long-term.

As punishment is local, it is to be expected that defectors would like to exist in groups that are far away from punishers, while punishers also prefer to not be in proximity to defectors as punishing incurs a cost to them. The patterns seen in Figure 17 mimic this. They also show that punishing and regular cooperators segregate, which supports the point made in Helbing *et al.* (2010) that punishing cooperators want to segregate in order to avoid being exploited by regular cooperators (who are second-order free-riding on punishing) and defectors. The emergent pattern of punishing cooperators outnumbering ordinary cooperators can be explained by the local nature of punishment: within a cluster of cooperators, a punisher may have the same utility as an ordinary cooperator and thus no incentive to change strategy. Further, punishers are effective at creating these clusters by disincentivising defection and thus may expand their cluster by converting defectors (something that regular cooperators maybe can not accomplish as well). Ergo, punishers may outnumber ordinary cooperators not because they have a higher utility, but because they are more effective at increasing their numbers by converting defectors.

Figure 17: Evolution of the SPGG with punishment (yellow) as an added strategy with $fine = 0.4$ and punishment cost = 0.1.




Observation 7: The SPGG simulation with punishment as an added strategy displays an average level of utility similar to the baseline, together with a trend of increasing utility over time.

Figure 18 displays the utility per capita over time for the punishment mechanism, and reveals that it is less volatile and has a similar average to the baseline when compared to Figure 15. The similar

levels could be explained by how utility gains from the increase in cooperation are counteracted by the utility losses from punishment fines and costs. There is also a trend of increasing utility over time, and there are three effects that likely contribute to this: (1) if the number of cooperators is increasing over time (which it does in this scenario), the public goods pot increases which benefits all; (2) when the number of defectors goes down and clusters of punishers have formed, the punishers no longer have a cost incurred to them since there are fewer defectors close to them; (3) like the second effect, when clusters of defectors have formed, it limits the total number of defectors that are exposed to punishers, so that only defectors on the boundary of the cluster are being punished. Thus, an increase in the number of cooperators increases the public goods pool and decreases total punishment, while the cluster formation also limits punishment by making sure that punishing interactions only take place on the boundary of the clusters, all of which progressively increases average utility over time.

Figure 18: Total group utility over time from the *Homo Moralis* SPGG simulation with punishment as an added strategy and with $\text{fine} = 0.4$ and $\text{punishment cost} = 0.1$.



5.4 SPGG with Reward Mechanism

Contrary to the ‘stick’ of punishment, reward mechanisms could prove to be useful ‘carrots’ for encouraging greater cooperation. In reality, different people can respond differently to punishments versus rewards, but there are most definitely scenarios in which rewards make sense. Consider efforts to encourage reticent individuals to get vaccinated in order to achieve herd immunity against an epidemic, thereby resolving a social health dilemma. If individuals are already rebelling against being told to get vaccinated by the government, being punished for not vaccinating might further embed their belief that they should not vaccinate, despite benefits to others in society from vaccination. For some of these individuals, rewards might be a better influence, and this is exactly what was seen in some countries during the covid-19 pandemic: some late-to-the-game vaccine recipients were coaxed into getting the jab by small rewards, such as lottery tickets or even direct cash payments. More generally, a society built on reciprocal rewards could go a long way in encouraging more cooperation.

Figure 19: Proportion of cooperation (ρ_c) over time in the SPGG with rewarding cooperation as an added strategy and with different rewards and costs of rewarding.





Table 7: Average proportion of cooperation (ρ_c) by specification of reward mechanism parameters.

Reward	Reward Cost	Average ρ_c ($T = 1000$)
0.2	0.1	46.87%
0.4	0.1	83.17%
0.2	0.05	61.95%
0.1	0.05	44.40%

Observation 8: Introducing rewarding cooperation as an added strategy has a positive effect on the proportion of cooperators in the game. However, the magnitude of the impact depends considerably on the strategy parameters, with the increase varying between 4-43 percentage points. In this scenario, the oscillations in the proportion of cooperators are largely suppressed, apart from one specification.

As with punishment, the average level of cooperation with rewarding cooperators is dependent on the mechanism's parameter specifications. However, the outcomes in proportion of cooperation (ρ_c) are more sensitive with the reward mechanism, as cooperation can be much higher but also lower than punishment, even though the costs and rewards are equal to the punishment fines and costs. The ratio between reward and cost is of particular importance, as observed in Table 7. A higher ratio (4:1 instead of 2:1) results in a much higher level of cooperation. But even when the ratio is the same, a higher absolute level of reward and cost seems to favour cooperation more, just like with punishment. This could be explained in a similar way, that is, when the rewards are higher in absolute terms, despite ratios being the same, fewer neighbouring rewarders around an agent are required for cooperation to give a larger payoff than defection. However, this effect seems to be more pronounced in the case of reward compared to punishment, perhaps because if rewarding cooperators all cluster around each other, they will all receive as well as give rewards, such that they end up with net positive returns to rewarding. Furthermore, one can see in Figure 19 that the dynamics of cooperation are more stable than the baseline, apart from the final case where reward is equal to 0.1 and the cost is 0.05, which has similar dynamics to the baseline with noisy fluctuations around the average. Overall, some significant benefits to cooperation are possible with the reward mechanism, such as an 83% average cooperation level, if such specifications are possible in reality.

Figure 20: Evolution of the SPGG with reward (green) as an added strategy and reward = 0.4 and cost = 0.1.




Observation 9: *The clustering effect seen in both the baseline and punishment scenario remains but now takes on a new form, as regular cooperators are drawn to the edges of the clusters of rewarding cooperators. Also, the amount of rewarding cooperators outnumbers the regular cooperators.*

As rewards are distributed locally by rewarding cooperators to all nearby cooperators, it is unsurprising to see clusters of cooperators and rewarders forming, as observed in Figure 20. Regular cooperators are pushed to the edges of these clusters because rewarders have an incentive to stick together, since when surrounded by rewarding cooperators they will get back more than the amount that they give to others, which they would not get from regular cooperators. Thus, regular cooperators still want to stick to the edges of these clusters to get some of the rewards from the rewarders, but are not generally “allowed in” to the centre. Also, since rewarders can actually have more utility within these clusters than regular cooperators who are not surrounded by rewarders, the amount of rewarders proliferates more than regular cooperators and ultimately outnumbers them.

Observation 10: *The SPGG simulation with rewarding cooperation as an added strategy displays an average level of utility that is roughly three times higher than both the baseline and the punishment scenario, albeit with a slightly declining trend over time.*

Despite the trend of slight declines over time, Figure 21 shows that utility per capita with the rewarding mechanism is on average around three times higher than in the punishment mechanism, when compared to Figure 18, and the baseline. This could be explained by a higher average level of cooperation, with this being achieved not by methods of costly punishment, but by a system of decentralised rewards in which the rewards are higher than the costs. Agents realise this and arrange themselves in clusters of reciprocity, and thus both cooperative and rewarding behaviour are proliferated across the game lattice, leading to overall gains in utility. This is an especially important finding, as it reveals in utility terms how much more useful rewarding can be than punishment if parameterised correctly, and also vindicates the importance of checking game outcomes not just in terms of cooperation levels and payoffs, but also in terms of utility.

Figure 21: *Total group utility over time from the Homo Moralis SPGG simulation with reward as an added strategy with reward = 0.4 and rewarding cost = 0.1.*



5.5 SPGG with Policing Mechanism

The policing mechanism bears much resemblance to the punishment mechanism, insofar as it introduces a ‘policing’ strategy through which cooperators can impose a punishment fine on defectors at a personal cost. However, the specific mechanism behind policing is slightly different, and as such the fine is referred to as a ‘policing fine’. Rather than being restricted to punishment of local neighbours, policing agents pay their policing cost into a global policing pot, which is scaled by the multiplier and then split across all defecting agents on the spatial lattice as a fine. This makes the mechanism more akin to a government service provision, just like a nation’s police force is funded out of taxes from all tax-paying individuals and is then used to encourage ‘cooperation’ with the law amongst the general public. It is therefore an interesting arena to explore the impact of government-provided services for encouraging greater cooperation, as one would expect an advantage because the global policing service also gains from the multiplier effect (as government services usually are ascribed a multiplier effect), which cheapens the cost of punishment for cooperating agents.

This mechanism reveals some basic dynamics as to how heterogeneously moral agents would approach paying for socially beneficial public services in the face of a social dilemma. To put this in real-world terms, imagine a nation-state within which individuals could choose what they pay tax for: in this state, a very stark social dilemma would present itself in areas like funding public policing services, which some individuals might wish to free-ride on at the expense of others. Of course, such a state is not very realistic, but more reasonably it is true that individuals are regularly asked to vote for a political leader amongst a pool of candidates who might offer different emphases on matters including police funding. Police funding was debated significantly in recent U.S. history, and the role of government in guiding public cooperation more generally through police and other similar forces has been a significant question in political philosophy for centuries.

Critically, it will be clear that the following mechanism is a very simplistic implementation of policing, and the results show an important limitation in the model design. However, this mechanism has been included to stimulate more creative thinking, as opposed to the more traditional notions of punishment and reward seen previously, and hopefully this could provide a foundation for more advanced assessments of public spending and other institutional incentives in areas of social dilemma.

***Observation 11:** Based on the first few rounds of the game, when there is a significant number of policing cooperators, it is clear that sustained policing would sustain nearly full cooperation.*

By the nature of the simulation design, agents are initially distributed on the spatial lattice at random, with equal likelihood of selecting one of the available strategies. As such, there is approximately one-third of agents choosing the policing strategy at the beginning. This creates a large pool of policing fines, made even larger by the multiplier effect connecting the policing cost to the policing fine, and this in turn means that defectors are punished significantly for their defection at the beginning of the game. Hence, the proportion of cooperation is driven to nearly 100% at the beginning of the game (see Figure 22), and it is evident that if a significant minority of agents could be encouraged to choose policing as a strategy in the long run – regardless of the magnitude of the fines/costs imposed – then a sustained level of high cooperation could be maintained.

Figure 22: Proportion of cooperation (ρ_c) over time with different fine sizes in the SPGG with a policing mechanism.




Table 8: Average proportion of cooperation (ρ_c) by specification of policing mechanism parameters.

<i>Fine</i>	<i>Cost (=Fine/Multiplier)</i>	<i>Average ρ_c ($T = 1000$)</i>
0.1	0.0625	46.98%
0.5	0.3125	48.77%
1.5	0.9375	48.50%
5	3.125	46.97%

Observation 12: In the context of this model, policing is a very unstable strategy at the beginning.


It is evident from Figure 22 that the high cooperation from policing is very unstable, and from Figure 23 one can see that policing drops to a low level after the initial higher proportion of policing. Taken without consideration of the model's limitations, this looks damning, as if policing is inherently unsustainable. However, it is important to recognise that the design of imitative strategy updating in this model is not well-suited for the policing mechanism: as backward-looking agents only learn from their payoffs and their immediate neighbours, once defection is almost entirely eliminated due to the large initial policing pool that punishes defectors heavily, there is no incentive to police anymore, because it is more costly than being a non-policing cooperator and now there are no

defectors left to police. Thus, policing becomes a weak strategy, and over a short period of time the defectors return. However, in reality one might expect commitment to a police force to be stickier, as forward-looking agents understand the need for policing as a *preventative* tool to stop defectors, much like Baumol *et al.*'s (1988) notion of the *threat of contestability*. That is, the mere presence of a meaningful defence against defection should discourage potential defectors from defecting in the first place. This is why the policing mechanism is a more abstract, illustrative mechanism to learn from, as opposed to punishment and reward.

Observation 13: *With policing as an added mechanism, high fines create a longer grace period of full cooperation at the beginning. However, low/medium fine levels appear to have a higher long-term average of cooperation.*

Unsurprisingly, given the damaging nature of high fines and the initial presence of many policing cooperators, high fine levels result in a more extended period of full cooperation at the beginning of the game, as this makes it much more painful to defect. Yet as with most other mechanisms, Table 8 reveals that it is low- to medium-size fines that are the most effective at increasing the average proportion of cooperation over the long-term of the game. As usual, this is driven by the lower cost of policing that is connected with lower fine levels, and in this case it is actually rather low costs that produce the strongest cooperation outcomes, because here there is a multiplier effect on the cost that produces a more significant policing fine, thereby deterring defectors at a lower cost. An implication here is that a low individual cost of policing could still achieve high cooperation if individual agents can be encouraged to contribute, because the multiplier attached to this cost produces a large global policing pot to punish defectors that can create high, if not full, cooperation.

Figure 23: *Evolution of the SPGG with policing (cyan) as an added strategy. In this specification, the policing $\text{fine} = 0.5$ for defectors and the cost of policing = $(\text{fine}/\text{multiplier}) = 0.3125$.*



5.6 SPGG with Reputation Mechanism

The first three mechanisms – punishment, reward, and policing – have revolved around monetary fines and rewards. However, as seen in much of the behavioural economics literature, nonmonetary mechanisms can be just as influential in shaping behaviour, if not more so (Kube *et al.*, 2012). As such, these next two mechanisms – reputation and moral adjustment – consider how agents respond to nonmonetary, social or intrinsic motivations to cooperate or defect.

To understand the reputation mechanism, consider a scenario in which individuals are asked in a visible, public setting to contribute to a charitable or other socially beneficial cause. Individually, people may not desire to donate, as it is costly and there is an incentive to free-ride in this social dilemma. However, having been asked in front of other people, there is a reputational element at play: many economic agents feel an acute social pressure to donate anyway, because they do not want to appear greedy or uncaring towards the charitable or social cause. Equally, agents may choose to donate specifically to look good in front of others. Thus, to protect or improve their social reputation, the agents may donate despite not doing so if it had been a hidden, private request. In the context of the SPGG, this means that agents still have only two strategic options – cooperate or defect – but there is a scalar reputation parameter on agents’ utility that increases every time agents cooperate, and decreases every time agents defect.

Observation 14: *With the added reputation mechanism, average cooperation is roughly 20-30 percentage points higher than the baseline scenario, depending on the specification.*


With reputation included as part of agents’ utility functions, there is an additional incentive for agents to cooperate, which translates into the substantial increase in the proportion of cooperators. Moreover, through imitative learning, agents begin to mimic the great performance of high-reputation cooperators in their immediate neighbourhood. As seen in Table 9, the range of average cooperation across multiple parameterisations is approximately 60-70%. Still, cooperation does not go to 100%, meaning that agents may still be in the presence of defecting neighbours whom they might imitate, taking a hit to their reputation but benefiting from the free-riding. Here, individual agents’ specific degree of morality (κ) in their *Homo Moralis* utility function will be particularly important in determining their likelihood to defect, as the intrinsic incentive to cooperate is stronger with reputation included.

Table 9: Average proportion of cooperation (ρ_c) by specification of reputation mechanism parameters.

Reputation Multiplier Range	Reputation Increment	Average ρ_c ($T = 1000$)
0.5 - 1.5	0.2	70.46%
0.5 - 1.5	0.1	63.94%
0 - 2	0.2	69.31%
0 - 2	0.1	64.49%

Observation 15: *There is an initial spike in cooperation, which then settles down and oscillates around an average level. The fluctuations are less frequent than in the baseline model.*

Figure 24: Proportion of cooperation (ρ_c) over time in the SPGG with an added reputation mechanism and with different specifications of the reputation multiplier and increments.



By observing Figure 24 and comparing to the previous examples, it is notable that the fluctuations around the average are lower than the baseline, but perhaps more substantial than some of the other mechanisms.

Observation 16: *The range of the reputation multiplier has no clear effect on the level of cooperation, but the reputation increment has a significant positive relationship with cooperation.*

As seen in Table 9, by comparing the average ρ_c between the cases where reputation is changed with an increment of 0.1 when agents cooperate or defect compared with where it is 0.2, it seems that the level of cooperation is much higher with a higher increment, and this effect is much more significant than that from increasing the range of the reputation multiplier.

Observation 17: *The clustering effect remains and spatial interactions are intensified with the increased cooperation from the reputation mechanism, as clusters appear to turn into ‘continents’ of cooperation/defection in various forms after $T = 10$.*

The ‘continent effect’ may be explained simply by the higher level of cooperation. It is also interesting that continuous defection in this mechanism leads to low levels of utility in the long run, so as seen in Figure 25, clusters of defection shift their location across the lattice over many rounds, since clusters of agents who defect continuously are susceptible to be wiped out by imitating cooperating neighbours who are doing better. Of course, another cluster of defection may subsequently appear elsewhere.

Observation 18: *The SPGG simulation with an added reputation mechanism has an average utility per capita that is almost double that in the baseline.*

The average utility per capita seems to be fluctuating in tandem with the overall level of cooperation as in the other cases, and the overall level of utility per capita is almost double that of the baseline, as Figure 26 shows. This is because the average level of cooperation is much higher and those who generally cooperate (which is most of the population) have a positive multiplier on their utility due to their higher reputation.

Figure 25: *Evolution of the SPGG with Homo Moralis agents over 1000 rounds with an added reputation mechanism to their utility, where the reputation range = $[0.5, 1.5]$ and the reputation increment = 0.2.*





Figure 26: Total group utility per capita over time from the *Homo Moralis* SPGG simulation with an added reputation mechanism, where the reputation range = $[0.5, 1.5]$ and the reputation increment = 0.2.



5.7 SPGG with Moral Adjustment Mechanism

This final mechanism is the second that deals with nonmonetary, intrinsic agent motivations. It is also the most novel mechanism presented here, having been created by the authors to consider how the *Homo Moralis* utility function and its κ morality parameter could be influenced over time. Moral adjustment here refers to the idea that agents' κ parameter may not be fixed over time; instead, it could evolve dynamically based on prior actions. That is to say, much in the same way that the reputation parameter increases for every cooperative decision and decreases for every decision to defect, in moral adjustment an act of cooperation will shift that agent's κ closer to 1 (pure morality), while an act of defection will shift it towards 0 (pure *Homo Oeconomicus* behaviour).

By way of example, consider how economic agents develop their moral preferences over time. The act of doing good things can make agents appreciate doing these good things more, thereby increasing their preferences for taking the morally 'right' choice, while agents who consistently do bad things might be adjusting their moral compass towards making more selfish decisions. Children learn a lot in this way: if their selfish behaviour is left unchecked, a child is likely to continue being selfish in the future. Moreover, children's moral awareness is greatly affected by learning from their surroundings. If their parents or other peers in their 'neighbourhood' display positive, morally conscious behaviour, they are more likely to grow up to prefer taking morally conscious decisions themselves, and they will get intrinsic utility benefits from acting in a morally 'right' fashion; whereas a child whose peers display selfish behaviour is likely to also defect from cooperative acts in social dilemmas, and not acquire utility from more cooperative acts. This is implemented in the model by a scaling factor attached to κ , which increases by a fixed increment every time a decision to cooperate is made, and decreases by the same fixed increment every time a decision is made to defect.

Observation 19: With moral adjustment as an added mechanism, the average level of cooperation is lower than the baseline (below 40%) apart from when the adjustment parameter is small.


There also seems to be a slight decrease in cooperation over time during the early part of the simulation in most scenarios.

As can be seen in Table 10, there is no considerable increase in the average proportion of cooperation from the baseline 40% level, and sometimes moral adjustment can lead to a slight decline in the level. This is because the main effect of moral adjustment on agents' decision-making is to alter their moral preferences in a way that pushes them further towards the extreme of their preferences; mildly selfish defectors evolve into strongly selfish defectors, while mildly moral cooperators evolve into strongly moral cooperators. Thus, the average proportion of cooperation does not change significantly over the course of the game.

Table 10: Average proportion of cooperation (ρ_c) by specification of moral adjustment mechanism parameter.

Morality Δ	Average ρ_c ($T = 1000$)
± 0.001	43.24%
± 0.01	39.26%
± 0.025	31.30%
± 0.05	36.54%

Figure 27: Proportion of cooperation (ρ_c) over time with different morality delta parameter sizes in the SPGG with a moral adjustment mechanism.



Observation 20: The main effect of moral adjustment appears to be that it makes the proportion of cooperation more volatile over time.

On a related note to the previous Observation, moral adjustment only makes agents' preferences more extreme. With more extreme defectors and cooperators, this can create a more volatile pattern of interaction between agents, as strong defectors clash with strong cooperators and individual utilities are driven to extremes, encouraging agents to try imitating their successful neighbours more. This volatility is visualised in Figure 27.

5.8 Robustness Tests

Given the path-dependent nature of agent-based simulations, robustness tests are critically important for ensuring results from the simulations are valid. In most cases, robustness tests will focus on the baseline model, but the final robustness test – which involves dynamising the public goods multiplier in various ways – is also applied to the reputation mechanism simulation, as this mechanism has been found to be the most reliably effective in encouraging a greater proportion of cooperation, and it is therefore relevant to check how robust this mechanism is to a more complex, dynamic multiplier environment.

5.8.1 Extending Rounds to $T = 10000$

Given the possibility of unpredictable emergent patterns and complex, chaotic dynamics, it is pertinent to check how robust the baseline simulation is to an extended time horizon. In this case, a horizon of 10000 rounds is chosen, as it reflects a 10x increase relative to the normal 1000 round limit, yet is still not too burdensome to simulate.

Observation 21: Over the longer time horizon of $T = 10000$, the average proportion of cooperation in the baseline model remains the same at 41%.

Figure 28 reveals that the baseline model's average proportion of cooperation is strikingly robust to an extended time horizon, coming in at 41% for both $T = 1000$ and $T = 10000$. This is encouraging to see, as it suggests that the game's long-term dynamics are relatively stable and inferences can be taken without significant sensitivity to the number of rounds. Similarly, Figure 29 displays similar clustering patterns to those in the normal-length game even at the later stages of $T = 2000$, $T = 5000$, and $T = 10000$. However, as with all economics models, and particularly with this inductive method of agent-based simulation, it is worth keeping in mind that structural shifts in the environment are not taken into consideration, so shorter-term dynamics are more reliably informative than dynamics that appear at a later stage of the simulation.

Figure 28: Proportion of cooperation (ρ_c) in the baseline simulation, with the number of rounds extended to $T = 10000$. Average cooperation across the game remains at 41%, as in the original baseline simulation.





Figure 29: Evolution of the SPGG with a distribution of heterogeneous Homo Moralis agents over 10000 rounds (blue agents are cooperators, while red agents are defectors).



Observation 22: *As in the normal-length game, the range of oscillations in ρ_c is stable between 30-50% throughout the extended-length game, and there are also similar intermittent periods of very low volatility in the ρ_c series.*

As was highlighted earlier in this thesis, one of the characteristics of complex adaptive systems like this SPGG environment is that while long-term equilibria are possible, unusual patterns can emerge from the system even if it is built purely on simple rules. Such patterns were identified in results from the baseline model, specifically the pattern of ρ_c volatility suddenly dropping to a very low level mid-way through the game. This finding also arises in the extended-length game, with a notable stretch of very low volatility in Figure 28 from approximately $T = 6000$ to $T = 7500$. However, the most important point to note here is that this remains consistent with the normal-length game, and Figure 29 shows how extended rounds still look very similar to the first 1000 rounds, so it appears that an extended time horizon poses no robustness concerns.

5.8.2 Varying Distribution of *Homo Moralis* Parameter

Observation 23: *Changing the distribution of the morality parameter, κ , in the agents' *Homo Moralis* utility function, fundamentally changes the baseline average of cooperation and fluctuations around the average.*

By comparing the results in Figures 30, 31 and Table 11 of changing the distribution of the morality parameter κ , one can observe as expected that this distribution is fundamental to the dynamics and results of the game. Therefore, correct specification of this parameter is crucial to the validity of the simulation. Overall, it is apparent that while some of these alternative distribution models can achieve a similar average level of cooperation across the entire game, it is the uniform distribution that best captures the style and magnitude of oscillations in the proportion of cooperation, as judged by experimental data from the public goods game.

Table 11: *Average proportion of cooperation (ρ_c) in the baseline SPGG for eight different distributions of κ , ranging from uniform and normal distributions ($U(\cdot)$ and $N(\cdot)$) to beta distributions ($\beta(\cdot)$) and gamma distributions ($\Gamma(\cdot)$).*

Distribution Method	Average ρ_c ($T = 1000$)
$\sim U(0,1)$	36.49%
$\sim U(0.4,0.8)$	43.09%
$\sim N(0.62,0.2), \kappa \in [0, 1]$	44.88%
$\sim N(0.5,0.25), \kappa \in [0, 1]$	31.60%
$\sim \beta(1.1,1.1)$	35.71%
$\sim \beta(3,3)$	29.40%
$\sim \Gamma(5,0.1), \kappa \in [0, 1]$	29.41%
$\sim \Gamma(2,0.2), \kappa \in [0, 1]$	21.10%

Figure 30: Evolution of the SPGG with varying distributions of the morality parameter, Kappa , in the Homo Moralis utility functions.





Figure 31: Further examples of evolution of the SPGG with varying distributions of the morality parameter, Kappa , in the Homo Moralis utility functions.



5.8.3 Varying Spatial Distribution of Agents at $t = 0$

Observation 24: *Changing the random number generator ('RNG') seed to vary the random distribution of initial agent location and all future random number generation within the code only changes the average proportion of cooperation (ρ_c) from the original 'Default' seed by a maximum of 1%.*

The slight differences in ρ_c seen in Table 12 might be explained by the fact that every complex adaptive system is path dependent, meaning that its starting characteristics can be meaningful dictators of how exactly the system moves. Likewise, this could arise from the fact that the RNG seed places agents of different morality into different locations with different types of agents surrounding them, and gives them potentially different starting strategies, making for a slight change in imitative behaviour. But since the overall effect is very small, changing the initial spatial distribution will not give significantly different results in terms of average cooperation over the course of 1000 rounds.


Table 12: *Changing the random number generator ('RNG') seed to vary the random distribution of initial agent location and all future random number generation within the code, and its effects on average proportion of cooperation (ρ_c). 'Default' seed is that used in the baseline model.*

Distribution Method	Average ρ_c ($T = 1000$)
Random (RNG seed = 'default')	40.93%
Random (RNG seed = 1)	39.94%
Random (RNG seed = 100)	40.23%
Random (RNG seed = 12345)	40.27%

Observation 25: *While the average proportion of cooperation across the entire game is robust to a changing RNG seed, the path-dependent nature of agent-based simulations makes the round-by-round pattern of cooperation very different.*

As already alluded to, this variance in the oscillation pattern seen in Figure 32 is driven by the fact that the specific movements within complex adaptive systems are determined by path dependency on initial conditions. Thus, the different patterns of randomisation in turn bring about different interactions between agents in the spatial grid and therefore different round-by-round imitative strategy updating patterns. Overall, though, the system maintains approximately the same long-term average proportion of cooperation, suggesting that there are no robustness concerns related to the specific style of randomisation.

Figure 32: Proportion of cooperation (ρ_c) over 1000 rounds of the SPGG simulation, for various random number generator (RNG) seeds used in the simulation code.



5.8.4 Varying Size of the Public Goods Multiplier

The public goods multiplier in this model is set at 1.6x to match the setup of Lugovskyy *et al.*'s (2017) experiment, whose results this thesis's baseline simulation aims to match for the first 15 rounds. Yet, as per Cartwright and Lovett (2014), the typical range of the multiplier varies between 1.3x-1.6x, and Zelmer's (2003) metastudy reveals that 1.3x, 1.4x, and 1.5x are all also commonly used values. This robustness test checks how resilient the model is to these multiplier values.

Observation 26: A lower public goods multiplier drives significantly lower levels of cooperation – from the 1.6x baseline multiplier, incremental reductions of 0.1x in turn reduce the average proportion of cooperation (ρ_c) by 5-7%.

Naturally, reducing the size of the multiplier does appear to reduce the incentive to cooperate (see Table 13), in line with expectations and results from meta-studies (Fiala and Suetens, 2017). As such, this robustness test also appears to align with expectations from the literature and theory, as it is obvious that lower multipliers should make cooperation less attractive. Thus, while the results

do indeed change for different values of the multiplier, consistency with the literature means that there are no robustness concerns here *per se*. Nevertheless, this highlights the importance of the multiplier value, and is suggestive that the size of the multiplier effect in real-world social dilemmas is an important factor to consider.

Table 13: Average proportion of cooperation (ρ_c) for different public goods multipliers in the range of 1.3x-1.6x. (Note: 1.6x is the multiplier used in the baseline model, and is included for comparative purposes.)

Multiplier	Average ρ_c ($T = 1000$)
1.3x	23%
1.4x	30%
1.5x	36%
1.6x	41%

5.8.5 Varying Parameter K in the Fermi Equation for Strategy Updating

In the baseline simulation, imitative strategy updating was determined by the Fermi-Dirac equation (Equation 8), within which K acts as a noise parameter affecting the probability of imitation. The baseline uses $K = 0.1$ as this is most common in the literature, but other values are also possible – for instance, Helbing *et al.* (2010) use $K = 0.5$. Thus, a range of possible K parameters are tested to see how robust the simulation is to this change.

Observation 27: Changing the size of K has no significant effect on average level of cooperation, but a great effect on volatility in cooperation.


Table 14: Average proportion of cooperation (ρ_c) for different values of K in the Fermi equation for imitation strategy updating. (Note: $K = 0.1$ is used in the baseline model, and is included for comparative purposes.)

K	Average ρ_c ($T = 1000$)
0.01	40.66%
0.1	40.93%
0.25	40.84%
0.5	40.65%

As seen in Table 14, the average proportion of cooperators does not change drastically with different levels of K . However, Figure 33 shows that the oscillations around the average level vary significantly based on the value of K . As a lower K represents less noise in the imitation probability (i.e. a lower K makes it more likely to imitate or stick with the strategy that will make them better off), it makes agents switch their strategy more often based on imitation. Nevertheless, the average proportion of cooperation does not even change by 0.5% across these specifications, and this is the most relevant result of the simulation, so there are no significant robustness concerns here. Nevertheless, it is wise to note that the shape and degree of volatility will be heavily affected by this parameter; sticking

with the standard value in the literature of $K = 0.1$ generates results that appear realistic when compared to experiments, so this might be an important parameter to commit to in further research.

Figure 33: Proportion of cooperation (ρ_c) over 1000 rounds of the SPGG simulation, for various values of the parameter K used in the Fermi equation.



5.8.6 Varying the Rate of Innovation in Strategy Updating

Imitation is the main way in which agents update their strategies, but innovation – that is, a random mutation of strategy that could be due to mistakes, a desire to try new strategies, or something else besides – is also permitted, albeit with a small probability of 0.1% per agent per round in the baseline simulation. This permits mutative behaviour change as is seen in the real world, and enables agents to potentially break out of a local equilibrium in their region of the spatial lattice. As such, the specified probability of innovation could be an important influence on game results, so is explored here in another robustness test.

Observation 28: Changing the rate of innovation has virtually no effect on the average level of cooperation, but a significant effect on the volatility around that average.


It is logical to expect that increasing the rate of random innovation in agents' strategies also increases

Table 15: Average proportion of cooperation (ρ_c) for different probabilities of strategy innovation updating. (Note: 0.1% is the probability used in the baseline model, and is included for comparative purposes.)

Probability of Strategy Innovation	Average ρ_c ($T = 1000$)
0.1%	40.93%
1%	41.20%
5%	41.56%
10%	41.88%

the volatility of strategy updating, and this is what the robustness test reveals in the charts in Figure 34. Given that innovation is by definition random here, it also represents a degree of noise, so again, more oscillatory patterns in simulation with a higher innovation rate are unsurprising. However, Table 15 reveals a very minor trend of increasing cooperation as innovation scales, suggesting that innovation does not play a large role in determining the average proportion of cooperation (ρ_c) in SPGG simulations, and no robustness concerns arise. Despite this, it is interesting that there appears to be a positive relationship between the innovation rate and ρ_c , one that is perhaps worthy of further investigation.

Figure 34: Proportion of cooperation (ρ_c) over 1000 rounds of the SPGG simulation, for various values of the probability of innovation-driven strategy updating.



5.8.7 Adding Semi-Stochastic Dynamism to the Public Goods Multiplier

The final robustness test is also perhaps the most important of this thesis, as it has a direct relevance to the research question. The research question asks not only how heterogeneously moral agents will behave in the SPGG in a more complex adaptive system; neither does it stop at asking how different social mechanisms can positively affect the proportion of cooperation. The final part of the research question goes as far as asking how robust the simulation and its social mechanisms are to dynamism in the public goods multiplier. This section assesses precisely this question, in the context of both the baseline simulation and the most successful social mechanism for improving cooperation, the reputation mechanism.

Four possible dynamics are created for the multiplier: first, it can simply follow a random walk, starting at the baseline 1.6x and changing randomly by increments of 0.01x each round; second, it can be a random walk with increments of 0.02x augmented with a small upward drift component of 0.00075x each round; third, it can be a cyclical multiplier following a sinusoidal curve pattern starting on an upward trajectory, and made semi-stochastic by adding random variation of 0.15x per round; and fourth, it can be the sinusoidal semi-stochastic cycle again, but starting on a downward trajectory.


Observation 29: *The proportion of cooperation (ρ_c) in the baseline simulation is, naturally, highly dependent on the specification of the multiplier dynamics.*

It has already been shown how the value of the multiplier is highly influential in determining the average proportion of cooperation. The results in Table 16 are therefore of little surprise – depending on the specific multiplier, the average ρ_c can vary significantly. Naturally, for example, the random walk with upwards drift has a higher average ρ_c because it is programmed to follow an increasing trend over time, thereby encouraging more cooperation. Notably, Figure 35 shows how the round-by-round cooperative dynamics are also very different across each simulation, broadly imitating the pattern of their multiplier dynamics. For instance, the average ρ_c for sinusoidally cyclical multipliers follow a similarly sinusoidal wave.

Table 16: *Average proportion of cooperation (ρ_c) in the baseline SPGG simulation for different patterns of public goods multipliers produced by semi-stochastic dynamic functions.*

<i>Function</i>	<i>Average ρ_c ($T = 1000$)</i>
Random walk starting at 1.6x, with 0.01x shocks	37.40%
Random walk with upward drift	50.78%
Cyclical with random shocks, starting upward	33.60%
Cyclical with random shocks, starting downward	33.31%

Figure 35: Proportion of cooperation (ρ_c) over 1000 rounds of the baseline SPGG simulation, for various patterns of public goods multipliers produced by semi-stochastic dynamic functions.




Observation 30: The proportion of cooperation (ρ_c) in the simulation with the reputation mechanism is also highly dependent on the specification of the multiplier dynamics, but it retains its outperformance and can in many cases reach full cooperation, if only temporarily.

Table 17 shows that, as in Table 16, the effects of different dynamic multipliers are powerful and varied. However, it appears that the better SPGG for cooperation – namely, that with the reputation mechanism – also performs better across the board in the face of these dynamic multipliers. This demonstrates how the relative performance of different social mechanisms against the baseline model are robust to more complex, dynamic multiplier situations, and moreover shows that reputation mechanisms remain a powerful tool in improving the average proportion of cooperation in spite of more complex or unforgiving multiplier patterns. Figure 36 reveals how in the presence of a random walk multiplier, average ρ_c remains well above 60% for most of the game, and in the event of random walks with an upward drift or cyclical multipliers, temporary or even sustained full cooperation (or almost full cooperation) is entirely reasonable.

Table 17: Average proportion of cooperation (ρ_c) in the SPGG simulation with an added reputation mechanism for different patterns of public goods multipliers produced by semi-stochastic dynamic functions.

Function	Average ρ_c ($T = 1000$)
Random walk starting at 1.6x, with 0.01x shocks	64.27%
Random walk with upward drift	90.35%
Cyclical with random shocks, starting upward	52.72%
Cyclical with random shocks, starting downward	51.22%

Figure 36: Proportion of cooperation (ρ_c) over 1000 rounds of the SPGG simulation with an added reputation mechanism, for various patterns of public goods multipliers produced by semi-stochastic dynamic functions.



6 Discussion

The novel nature of this thesis and its broad scope of findings give rise to many areas of discussion. For clarity, and to highlight this thesis's different areas of contributions, these are split into distinct categories as follows.

6.1 Overall Trends

Before diving into specifics, four overall trends have been identified that are worthy of being highlighted. These four points alone offer useful insights for economists and policymakers as to how to better understand cooperation in social dilemmas, particularly when taking into account the uncertain nature of complex adaptive systems like those in which economic models and policy decisions must be made:

- (i) ***The agent-based SPGG simulation offers both breadth and depth in exploring economic behaviour:*** Through one relatively simple agent-based simulation, it has been possible to reveal both the limitations of the analytical game-theoretic method in a complex adaptive system (namely, the total lack of cooperation when using only *Homo Oeconomicus* rational agents) and the power of different social mechanisms to increase the average proportion of cooperation (ρ_c) in the spatial public goods game. Furthermore, the simulation method has proven not only to be easily calibrated to mimic experimental results in the baseline model, but also to be robust in the face of different initial conditions and other changes to parameters. All of these benefits combined enable the authors not only to explore a wide range of factors influencing and encouraging cooperation in the context of complex adaptive systems, but also to delve deep into how different factors might affect cooperation and how various types of complexity may influence results.
- (ii) ***Policymakers should consider their 'Goldilocks' variable ranges for mechanisms to improve cooperation:*** In a majority of social mechanisms, it is neither the largest nor the smallest fine/reward/reputation range that encourages the highest average proportion of cooperation (ρ_c); rather, it is a specification of medium magnitude. This is reminiscent of the 'Goldilocks' notion of 'just-right' parameters for success, and it is surprisingly robust across the different mechanisms. For punishment and reward this is less clear, but in these mechanisms only low-to medium-sized fines were explored, so they neither support nor contradict this position. This corroborates similar findings in both the behavioural economics and SPGG simulation literature, and moreover it does so with more rigorous economic structures and dynamics like utility functions and heterogeneous behavioural preferences. As such, policymakers may benefit in future from selecting mechanism parameter values of medium magnitudes, in order to obtain the highest possible levels of cooperation in social dilemmas.
- (iii) ***Intrinsic, nonmonetary mechanisms can be even more effective than extrinsic, monetary mechanisms:*** Corroborating inferences in the behavioural economics literature (e.g. Masclet *et al.*, 2003), it has been found that nonmonetary mechanisms like reputation that have more to do with an economic agent's intrinsic belief system than monetary rewards or punishments are more reliably influential in encouraging greater levels of cooperation. This can be due to social image, that is, a fear of looking bad in front of others or a desire to 'show off' how good one is

by opting for a cooperative strategy. Policymakers may be able to make use of this by designing policy in social dilemmas to appeal to agents' social reputations, thereby encouraging the most cooperation whilst also not having to spend a lot on managing punishment/reward systems or penalising cooperators by the cost of punishing or rewarding.

- (iv) *Utility is an important extra dimension for assessing outcomes, revealing a potential benefit of rewarding and an issue with punishment:* It has been found in this thesis that reputation is the most consistently beneficial social mechanism for improving cooperation, and it also roughly doubles average per capita utility during the game. However, comparing utility outcomes for each social mechanism, it becomes clear that rewards can actually produce average per capita utility that is three times greater than in the baseline, because the virtuous cycle of rewarding cooperators giving rewards to each other can result in all of them being better off, as the reward they receive is always greater than the cost they pay to reward others. Meanwhile, punishment results in a similar average per capita utility than in the baseline, because the utility benefits of more cooperators are counterbalanced by the extra disutility from punishments and their costs. Utility therefore offers an additional dimension for reviewing mechanisms' effectiveness, which is not usually considered in the SPGG literature.

6.2 Homo Moralis and Bounded Rationality in the Spatial Public Goods Game

Assuming perfect rationality, according to standard economic theory, agents playing the spatial public goods game should realise that defection is a strictly dominant strategy over cooperation, since an agent's payoff, and subsequently their utility, is always higher when they defect. Moreover, even when restricting agents to a bounded rationality where they cannot perfectly compute the dominant strategies in a complex environment and are instead forced to compare with and imitate their neighbours, when the agents' utility functions are of the *Homo Oeconomicus* variety, agents updating their strategy via imitation and innovation still arrive at a zero-cooperation outcome. Thus, a different utility function needs to be supplied.

Initially, the authors tried popular utility functions from the behavioural economics literature such as Fehr-Schmidt Inequity Aversion and Charness-Rabin Conditional Welfare, but replaying the game with these utility function still resulted in the social dilemma and a steadfast state of non-cooperation, as shown in Appendix A. As discussed in the appendix, this is partially due to the definition of these utilities not just punishing defectors who do better, but also giving cooperators with a low payoff some incentives to defect. The main reason, however, might be that these utilities are strictly based on monetary payoff, something that the *Homo Moralis* utility function avoids.

As long as agents' decisions to cooperate or defect are solely based on payoffs, without forceful added mechanisms to change the economic incentives, it is likely that these simulated public goods games result in near-zero cooperation due to the inequality in short-term monetary payoff between the two strategies. So to create a baseline simulation that resembles real-life experiments made on the public goods game (which do not have added mechanisms), a utility function that was not only based on monetary payoff was required. The advantage of *Homo Moralis* is precisely this, as agents are given a heterogeneous preference for doing the 'right thing' by adding an argument in the utility equation that values what the outcome would be if everyone chose the same strategy as them. As an outcome with full cooperation brings greater utility than a state of full defection, agents experience an internal reward for choosing to cooperate, and results mimicking experimental data can finally

be obtained.

This further adds to the realism of the simulation, since the decision to cooperate or defect in a setting like the spatial public goods game is not merely an economic choice, but also a moral one. This simple – and underappreciated – distinction could explain why some economists are bewildered by the existence of ‘irrational’ but prosocial economic behaviours that are present in society, while other fields of researchers are not. Further, it is an added benefit that the bounded rationality defined in this paper (imitation- and innovation-based strategy updating rules) combine well with *Homo Moralis* to create complex patterns and dynamics of cooperation that are comparable to the findings in the experimental literature.

6.3 Social Mechanisms in the Spatial Public Goods Game

6.3.1 Reputation

Of the mechanisms that have been proposed in this thesis, reputation arguably has the most reliably effective impact on cooperation and utility levels across different specifications. This is supported by the fact that the average proportion of cooperators is increased by between approximately 25-30 percentage points in all scenarios. A reasonable counterpoint is to question whether this efficacy is coerced by the specification of the reputation parameters, as a multiplier on utility ranging from 0 to 2 based on reputation can seem a bit extreme. One reply would be that the range of this multiplier does not seem to have that much of an effect on cooperation as seen in the results, and that the most important parameter is the increment by which an agent’s reputation is changed by their choice of strategy. An increase in reputation by 0.2 or 0.1 points does not seem particularly extreme but makes a significant difference in the behaviour of agents playing the game.

The relative effectiveness of the reputation increment parameter stems from the strategy updating process. Agents’ choice of strategy depends on the difference in utility between their last round and what utility they would have had if they selected the strategy of their ‘best performing’ neighbour, in other words the one with the highest utility. As they assess this difference, they take into account how much their reputation changes short-term by selecting the other strategy. As agents only consider short-term effects, this reputation increment can have a significant impact on their choice even if the increment is small. Hence, by assessing the short-term impact of their strategy choice on their reputation, and the reputation parameter being a scalar on their overall utility, agents are incentivised to cooperate more than previously.

The mechanism also synergises well with the heterogeneous distribution of morality in the *Homo Moralis* utility function, as the two components have a few interesting interactions. For example, agents that are more concerned about their monetary payoff may be more inclined to suffer a loss of reputation for the monetary gain, and moral agents are given additional rewards for their preference to cooperate. The mechanism also adds its own additional characteristics to the game for agents who are at the boundaries of reputation values. Usually, the net difference in reputation by choosing one strategy over the other is double the reputation increment, since if an agent chose the other strategy their reputation would move an equal increment in the other direction. However, for agents who have reached either the minimum or maximum of the reputation interval, the net difference is halved as their reputation may only move in one direction, or stand still. Thus, an agent who has

continually defected over multiple rounds and finds its reputation to be at the bottom level, will lose less utility for defecting than other agents, while habitual cooperators also gain less utility for cooperating. Therefore, even if the mechanism still incentivises cooperation, the incentives are half as effective on the boundaries of reputation.

It is important to note that this mechanism is not an economic incentive, but an internal process of agents being concerned about their morality or self-image. Reputation is a ‘natural’ mechanism, since there is a pronounced evolutionary advantage associated with this feature of human behaviour, and it does not rely on monetary outcomes. It is a social mechanism that can be observed throughout all of society, not least in group projects where free-riding dilemmas may arise. Therefore it is also no surprise that it has the potential to be effective in promoting cooperation in the spatial public goods game, although this is admittedly somewhat dependent on its mathematical definition. In conclusion, the reputation mechanism gives agents a further internalised ‘taste’ of morality, whether egocentric or not, in addition to the *Homo Moralis* utility. As a concern for reputation is an intrinsic and costless feature of individuals while also being a social mechanism not based on monetary outcomes, intensifying the importance of reputation in a social setting could be an effective method of promoting cooperation when the incentive to defect is large, like in the spatial public goods game.

6.3.2 Punishment and Reward

As punishment and reward are similarly structured – yet antithetical – mechanisms which share the same parameter values, some direct comparisons of their effectiveness and characteristics can be made. The main differences in characteristics between punishment and reward is that for punishment, ‘everyone loses’ in terms of utility unless there is an increase in cooperation that outweighs the loss both by increasing the public pot and decreasing punishment. However, with reward the giver loses utility, but other cooperators benefit. If agents give to each other then it would be mutually beneficial, and they would end up with more in the end because of the reward being higher than the cost. Further, due to the local nature of the strategies, both mechanisms display dynamic spatial patterns, such as clustering and segregation. Punishers tend to cluster with other punishers to lessen their cost of punishing and to avoid ordinary cooperators who free-ride on punishment. They also incentivise neighbouring defectors to convert to a cooperative strategy, which becomes a method for punishing clusters to expand. Rewarders also cluster with other rewarders for similar reasons, since they avoid ordinary cooperators who incur a cost and no benefit to them.

In terms of effectiveness, punishment is a social mechanism that may be aptly described as a double-edged sword regarding its social impact. It has a rather substantial and robust positive effect on average levels of cooperation in the spatial public goods game. As can be observed in results, the effect on average levels of cooperation varied between an increase of roughly 8-20 percentage points, depending on the specification. However, even in the case when cooperation increased the most, the level of average utility remained the same as in the baseline without the mechanism. Thus, any gain in well-being from increased cooperation is offset by the punishments and the costs. As the purpose of this thesis is to investigate effectiveness in promoting cooperation and its broader context of social utility, punishment does classify as an effective social mechanism since the mechanism was able to increase cooperation a significant amount across different parameter values. However, the

cost of this increase in terms of utility is an important side-effect to consider.

Reward on the other hand can have a substantial positive effect on both cooperation and social utility, however it is not as robust in its effects as punishment, as the effectiveness varies widely between different parameter values. In two out of four specifications, average cooperation only increases by 4-6 percentage points, but in the best-performing case the average proportion of cooperators increases by 43 percentage points, which is a doubling in the level of cooperation. In this scenario, average utility is also around three times higher than both the baseline and punishment. Robustness in performance across specifications is a good indicator of an effective mechanism, however the best-performing cases are important to analyse as well. The classification of reward as an effective strategy thus depends on whether the parameter values that caused the large positive effects are realistic, which might be impossible to answer definitively. Thus, it is fair to say that the reward mechanism has significant potential to promote cooperation in a social setting like the spatial public goods game, although it is also likely to not have a large impact depending on the parameters. Neither of these mechanisms however, are as consistent at promoting cooperation across specifications as reputation.

Some potential pitfalls when defining these mechanisms is to not put a limit on the costs that an agent can incur to itself when rewarding or punishing another agent. Some of the literature sets very high fines and rewards which necessitates high costs that the agents may not in fact be able to pay due to the limited endowment given to the agents at the start of each round. Thus, as an extension to this thesis, it could be possible to build a version with credit markets where agents can borrow to punish or reward beyond the size of their endowment. However, this thesis's current model remains robust since agents suffer negative utilities from spending beyond their means, and also because credit markets are absent in every other spatial public goods game paper that has been published so far. Therefore it is not a substantial factor, but agents' financial means are nevertheless worth considering when specifying these mechanisms.

6.3.3 Policing

As a somewhat less important, but still meaningful social mechanism for improving cooperation, two things ought to be discussed about policing. First of all is to recognise the limitations of policing in this thesis – as already mentioned, the simple strategy updating used here to mirror the literature's main approach of modelling means that once the broad success of policing wipes out defectors, the incentive to choose the policing strategy drops away, and agents never learn effectively to return to this strategy despite its clear long-term benefits of dissuading defection. If, instead, a small but mandatory policing requirement was imposed upon all agents, it is clear how that could lead to sustained and nearly perfect cooperation in the context of this model. Ergo, the policing mechanism could pose significant benefits to policymakers in stimulating cooperation if implemented in a tactical way like this.

Second, much debate was had between the authors over how exactly to specify the parameters for policing. Setting very high policing fines alongside low costs was deemed to be 'begging the question', a logical fallacy that would not help to make this thesis robust. At the same time, it is plausible that one could indeed model a very high policing fine with very little cost. For example, successful policing in the real world can result in hefty prison sentences, or even manual labour or capital punishment in some cultures, which can be seen by the agent as substantial hit to their

utility at a limited cost to non-defectors. Thus, subsequent research might do well to explore how exactly different levels of policing fines and costs can improve cooperation, and to examine the feasibility of these specifications in the real world.

6.3.4 Moral Adjustment

The final mechanism to discuss is moral adjustment. This does not require too much focus, as it is the most novel of all mechanisms considered in this thesis: the concept has been devised from scratch by the authors, as a way to investigate how the *Homo Moralis* morality parameter κ might change over time. However, its novelty is also its major limitation, as there is no background literature on this sort of mechanism.

At a high level, the mechanism does offer some interesting ideas, as it offers a new way to interpret how agents might see their moral preferences change dynamically over time. This challenges some models of altruistic behaviour in behavioural economics that tend to assume that agents have fixed preferences regarding inequity aversion, altruism, and so forth. However, the mechanism is shown only to create more volatile cooperation series, with little impact on average cooperation, so this may present supporting evidence that ‘fixed’ parameters might be reasonable over the long term, across entire groups at least. No concrete conclusions can be made from this basic application of a moral adjustment mechanism, however, and more research into how moral preferences might change and how best to model this ought to be conducted before applying these observations to real-world situations.

6.4 The Public Goods Multiplier

Retesting the baseline and reputation SPGG simulations with dynamically evolving multipliers is more than just a robustness test; it is also analogous to how real-world multiplier mechanisms can operate. Take, for example, financial markets: participants know that markets will go up in the long-term, but in the short-term there is a large degree of uncertainty regarding how positive gains will be, if positive at all. This is emulated by one of the dynamic multiplier patterns chosen in this test, the random walk with upward drift. More generally, financial markets and the broader real economy follow broadly cyclical patterns, as represented by the sine wave variations of the dynamic multiplier used. Yet these sorts of fluctuations in potential payouts are not considered in the vast majority of analytical game theory and behavioural experiments, so the perceived ‘best’ strategies and ‘right’ behavioural trends found in these studies might not actually be the most relevant in more realistic settings of dynamic uncertainty.

Of course, the specific programming of these dynamic multipliers will weigh into how exactly they affect different SPGG simulations, as will the level of nuance allowed in agents’ strategy updating rules and utility functions. Thus, there are certain limitations in inferring real-world behaviour from these simulations. Nevertheless, this robustness test of dynamic multipliers appears to offer deeper insights into the value of the reputation social mechanism, which holds up remarkably well in each of the multiplier patterns. Going forward, subjecting the findings of analytical game theory and behavioural experiments to this sort of robustness test could help to test the external validity of these findings, and ultimately enable policymakers to have more confidence that they will hold

up in the face of real-world uncertainties.

6.5 Robustness in Simulating the SPGG

Observations from the various robustness tests all align with one core point: the SPGG simulation created in this thesis is surprisingly resilient to changes in key parameters and model design choices. One piece of early feedback noted in the literature and received from simulation experts was that simulations are often heavily dependent on their starting parameters, and path dependence from these initial conditions can lead to wildly different outcomes in the game. Yet this thesis has experienced no such issues, despite conducting a wide range of robustness tests. Naturally, varying some initial parameters like the magnitude of the public goods multiplier has a noticeable effect on results, but this is to be expected as values like the multiplier are crucial in determining agent payouts, which in turn directly influences utility and strategy updating via imitation. Moreover, the multiplier was selected to mirror a common value in the literature, so it is in line with the rest of the research field. In sum, the methodological approach of agent-based simulations taken by this thesis has proven to be surprisingly robust.

6.6 Limitations

While the SPGG simulations in this thesis have been built on a distinct branch of the literature, and robustness tests have all checked out in a satisfactory manner, there are of course limitations. One category of limitations is to do with the choice of method itself. This has already been discussed in Section 3, but briefly, the agent-based simulation method lacks the logical certainty that can be provided by analytical, deductive methods. It also can be seen as dependent on modellers' specific design choices. The latter problem has been mitigated as far as possible by (i) relying on the literature for specifying the vast majority of the simulation's parameters, and (ii) placing great emphasis on robustness tests to defend against the possibility of modelling choices producing unreasonable outlier results. Yet the former problem is harder to combat directly: there is no way to get around the fact that inductive methods like simulations provide less certainty than deductive methods. On the other hand, one can recognise that while simulations offer less certainty, they also rest on less restrictive assumptions than analytical methods, so increase realism and arguably can be designed to provide more detailed insights into shorter-term system dynamics. Thus, if complex systems are of interest, the agent-based simulation method may still be seen as most attractive.

Focusing more specifically on aspects of the simulation design, other limitations ought to be highlighted. It is impossible to fully capture the complexity of human decision-making or real-world uncertainties, but some simulation design elements must be set. This thesis has tried, as far as is possible, to mirror the design of other SPGG simulations and public goods experiments as possible. As such, the vast majority of parameters have been set based on those used in the literature. One critical area of divergence in this thesis has been the introduction of utility functions, particularly the *Homo Moralis* utility function. To specify this, a decision was taken to give all modelled agents a heterogeneous value for their morality parameter $\kappa \in [0, 1]$, drawn from a uniform distribution between approximately 0.1 and 1. This not only accurately calibrates the first 15 rounds of the simulation to match round-by-round results from a real public goods experiment, but it also closely matches the proposed range of parameter values given by the creators of the *Homo Moralis* utility

function (Weibull *et al.*, 2020), and it survives robustness tests in Section 5.8.2. Nevertheless, this section does depart from the literature, so deserves the most scrutiny by future researchers keen to pursue a similar simulation approach.

Finally, an area of model design which possesses limitations despite it being taken directly from the literature is the style of strategy updating rules for agents. A simple strategy updating rule of imitation dominates how agents update their strategies each round, based on what they observe their immediate neighbours doing. However, this can be seen as too simplistic, given the manifold drivers that are often behind why economic agents really make decisions. Various alternatives could have been pursued here, from the random inductive rules used by *e.g.* Arthur (1994), which could include “switch strategy every 4 rounds” or “copy the highest-utility neighbour every second round”, to other game-theoretic strategies like tit-for-tat or the grim trigger strategy. Equally, agents could have been given a memory of their and their neighbours’ performance over the past few rounds, to devise a more holistic picture of which strategies might be best for them; or instead of only backwards-looking inductive reasoning, agents could be at least partially endowed with an ability to look ahead to see the longer-term benefits of their strategy choice, rather than only looking backwards in time for information to guide strategy updating. This thesis chose the simple imitation strategy because it already innovates in its general method in a few other key places, such as the introduction of utility functions with moral preferences, so it was deemed sensible to mirror the literature’s approach elsewhere. But if there is an obvious area for further exploration it is here. Overall, this thesis has had to make some difficult decisions to copy the literature’s approach and pursue more economically meaningful approaches, and the authors hope that while it lays a strong foundation for SPGG simulations, further refinement and exploration will be pursued by other researchers in the future.

7 Conclusion

This thesis has sought to better understand cooperative dynamics in social dilemmas, with a research question focused on better understanding how to improve cooperation with various social mechanisms available to economists and policymakers, and to test the robustness of these cooperative dynamics and social mechanisms in the context of complex adaptive systems with dynamic, unpredictably evolving system parameters. This thesis has addressed the cooperative social dilemma by studying the canonical public goods game in the behavioural economics literature, and it has done so by applying the method of agent-based simulations to a spatial version of the public goods game. Most design features of the simulation, as well as specific parameter values, were taken from the interdisciplinary literature so as to ensure relevance, reliability, and consistency with the research field, although steps were taken to increase the amount of economic content by adding features like utility functions and bounded rationality. With just one free parameter left to determine after this, care was taken to calibrate the simulation in such a way that even allowing for heterogeneous preferences, cooperative dynamics in the baseline model closely mimic the amount of cooperation round-by-round in real public goods game experiments, thereby offering a more realistic environment for exploring how social mechanisms and dynamic uncertainty influence the level of cooperation during the game. Testing of five different social mechanisms – punishment, reward, policing, reputation, and a self-defined ‘moral adjustment’ mechanism – combined with a wide collection of robustness tests enabled this thesis to address both the benefits of each social mechanism, and the realism and robustness of the complex social system that was created.

The purpose of this thesis was effectively split into three intended outcomes, as explained in Section 1.1. Addressing each of these in turn:

- (i) ***To better understand how economic agents cooperate in a realistic, complex world:*** The results gathered through the agent-based simulation method of the spatial public goods game have directly contributed to a better understanding of how economic agents cooperate in a realistic, complex world, both with and without social mechanisms to assist cooperation despite the social dilemma. Reputation was found to be the most reliably useful mechanism for improving cooperation, while punishment and reward systems were also successful in improving cooperation, if a bit more volatile based on the precise specifications of their fines, rewards, and costs of fining/rewarding. The complexity of the system was achieved by granting agents heterogeneous preferences as well as bounded rationality via a strategy updating rule based on the performance of immediate neighbours on the spatial grid, and was further tested through robustness tests.
- (ii) ***To more clearly establish the economic content of agent-based simulations:*** Agent-based simulations of the public goods game, which to date have been predominantly researched in an interdisciplinary fashion by evolutionary biologists, complex systems theorists, and physicists, have been somewhat lacking a few core economic concepts. This thesis crucially added the notion of utility to replace simple monetary payoffs or the notion of biological ‘fitness’; enabled a richer variety of social behaviour by enabling heterogeneous moral preferences; and tailored some social mechanisms to better reflect parameterisations and theories in the economics literature. The relative success with which the results were able to emulate real experimental data in the initial calibration while taking virtually all parameters from the relevant literature, on top of the robustness of the simulation to various changes and extensions to initial conditions

and dynamically evolving parameters, is suggestive that the additional economic content included in the SPGG simulation of this thesis helped considerably in aligning the standard SPGG literature with economic theory and reality.

- (iii) ***To more clearly establish the complementary role that agent-based simulations can play in supporting the external validity of analytical and experimental findings in economic research:*** By testing the spatial public goods game with the agent-based simulation method, it was possible to subject findings from analytical game theory and behavioural mechanisms to external validity checks in a new, complex setting. For example, the usefulness of punishment, reward, and reputation systems – which have obtained much attention in the literature – were assessed not only in the baseline simulation, which already includes a significant degree of complexity in the system and bounded rationality in the agents, but also settings with even more complex, dynamically evolving parameters. The results have broadly corroborated findings in the behavioural literature as well as in the SPGG simulation literature, providing them external validity, while adding some nuanced evidence to debates on e.g. the limited applicability of canonical behavioural utility functions like the inequity aversion model (Fehr and Schmidt, 1999) or the conditional welfare model (Charness and Rabin, 2002), and the potentially more viable *Homo Moralis* utility function of Alger and Weibull (2013). Finally, the method has also helped to detect novel emergent patterns and trends that analytical and experimental methods might otherwise miss, such as the importance of setting social mechanisms with parameters of medium magnitudes, as opposed to e.g. very large or very small fines, and also the tendency of the proportion of cooperation to fluctuate around a long-term average with fairly uncertain shifts between low- and high-volatility phases.

Overall, this means that not only has the research question of understanding how social mechanisms can improve cooperation in a realistic, complex setting been answered, but the remaining goals of increasing the economic content of this currently interdisciplinary topic and establishing a secondary role of agent-based simulations as a way to test the external validity of more traditional research have also arguably been achieved. This presents many fruitful applications and avenues for further research, as will be discussed in the final two sections.

7.1 Potential Applications

Topic-wise, the broad applicability of the concept of cooperation in social dilemmas means that there are many topical areas for consideration. *Conflict* or *war strategy* is one broad area of application – Thomas Schelling won his Nobel prize for applying basic social simulations to the topic of conflict, and game theory has long been dealing with these questions, so the ability to add significant systemic uncertainties and complexity at fast pace and low cost makes agent-based simulations a ripe method for assessing and adapting various war and other conflict strategies to better understand how to achieve cooperative outcomes in real-world settings. *Climate change* is another obvious and important topic in the international policy arena, as while global cooperation is required to address climate change mitigation and adaptation, there is a significant degree of defection from global agreements by individual countries and other meaningful actors. Even *public health*, and specifically *epidemiology*, could be an interesting area for further explanation, with the spatial social dilemma setting being relevant for vaccination games as discussed in Section 2.4. More generally, group cooperation in

social dilemmas is often relevant in the areas of *public taxation and public goods provision*, *firm decision-making*, *management techniques*, and much more. This list is surely not exhaustive, but highlights the flexibility of the topic and also of the simulation method.

Considering ‘applications’ in an even broader sense, one can consider the different purposes agent-based simulations of social dilemmas and other games can be used for. Two purposes have been highlighted here: first, for exploring cooperative dynamics and emergent patterns from complex economic systems; and second, for testing the external validity of findings from other, more traditional methods. In this way, agent-based simulations can be seen not as a rivalrous method to those used previously, but rather as a complementary method to the current epistemological toolkit of economists.

7.2 Areas for Future Research

As this thesis has primarily sought to unify the research paradigms of complexity economics and agent-based simulations with more traditional economic and behavioural theory, it has only focused on a small subset of possible avenues to explore. Furthermore, the vast research space offered by complex systems, as well as the highly flexible nature of agent-based simulations, open up both broad and deep areas to investigate. From this thesis, there remain at least a few valuable areas for further research:

- **Enhancing Adaptive Learning:** The agents in this model are entirely myopic, and only consider what they have learnt in the most recent round. However, it is perhaps more reasonable to permit agents to have a longer memory of, say, up to 3 rounds, and to update strategies based on this greater set of information. In such a scenario, the notion of hyperbolic discounting of future utility could also be added to imbue agents’ recollection of payoffs with time inconsistency, further integrating behavioural economics content into the model. Finally, elements of communication between non-neighbouring agents via random cross-grid connections or imperfect communication across a wider set of neighbours could enrich the set of influencers for each agent and more closely approximate the more globalised world of today.
- **Alternative Strategy Updating Rules:** Beyond the adaptive learning strategy updating rule used in this model, there remain a wide array of alternative possible inductive reasoning techniques that agents could realistically use. Rather than simply learning from agents around them, it could be interesting to investigate how agents might retrospectively evaluate what would have been their best possible strategy in each round (the *best response* updating rule), or to imbue agents with traditional strategies like *tit-for-tat*, *grim trigger*, or *semi-grim* (although these are more difficult to implement in the public goods game specifically). In a different setting, Arthur (1994) gave agents up to 30 simple rules, akin to “defect every second round” or “follow the best strategy from the average of the past 4 rounds”. Memory of the past few rounds of personal outcomes and neighbours’ strategies could also be built into a more advanced strategy updating rule to better reflect how agents make economic decisions in reality (Bordalo *et al.*, 2017).
- **Exploring the Social Network Structure and Information-Sharing Mechanisms:** As reported by Jackson *et al.* (2016), the initial setup of the social network (in this case, locations of agents

on the square lattice) can have meaningful effects on how they learn and evolve over time. However, they also note that a high degree of similarity in local neighbourhoods does not necessarily create impermeable clusters; rather, this can help to foster new strategies or innovations that would not occur otherwise. Given its similarity to the typical procedure of a randomised laboratory experiment, this thesis's choice to randomly scramble agents' positions on the spatial lattice was necessary to be able to mimic experimental data, but further investigation of the role of the structure of social networks could offer deeper insights into how economic agents interact in reality. Similarly, further explorations could assess how agents update their strategies if they receive only incomplete information about their neighbours' actions or the game parameters (De Marti and Zenou, 2015).

- **Exploring Local vs. Global Punishment/Reward/Policing Mechanisms:** This thesis has decided to implement punishment and reward at the local level, based on the level of defection in each agent's immediate neighbourhood, while policing has been implemented at the global scale of the entire game. Further exploration of the differences in cooperation effects between local and global implementation, and other related nuances like the relative costs and fines that can plausibly be associated with each of these scales of implementation, could offer very practical insights for policymakers.
- **Adding Credit Markets for More Accurate Budget Constraints:** One area of relaxed assumptions in this thesis – as well as in the SPGG literature more broadly – is budget constraints. That is, while agents are modelled as receiving an endowment of 1 in each round, they can contribute their 1 and still pay additional costs for punishment or reward, or even enter negative endowment territory if punished at a high level. This could be adjusted by considering a limited credit market permitting agents to borrow at additional cost if they enter negative endowment territory, and limiting the degree to which they can become or remain negative. However, this was not deemed to be the most important contribution in the context of this thesis, so the literature convention of ignoring such issues was followed in this case.
- **Alternative Games:** The prisoner's dilemma is the most explored behavioural game, and the public goods game has also been researched a great deal. Yet there are many other games that could be simulated, such as the *dictator game* (e.g. Dana *et al.*, 2006) or the *trust game* (e.g. Cox, 2004), which could realise more insights from complex environments and dynamics.

Moreover, this subset of behavioural economics is just one of many fields of economics that agent-based simulations can be applied to, with significant potential in areas such as financial markets and macroeconomics. With so much left to explore, there is hence a significant degree of future work left for other researchers to pursue. So, in answer to the question of “what comes next?” at the 2021 Nobel Symposium “100 Years of Game Theory” mentioned in Section 1, it appears that even if research avenues of analytical game theory are close to exhaustion, the possibilities arising from agent-based simulations in economics have only just begun to be explored.

8 References

- Abdou, M., Hamill, L., and Gilbert, N. (2012). Designing and Building an Agent-Based Model. In: Heppenstall, A. J., Crooks, A. T., See, L. M., and Batty, M. (Eds.). *Agent-Based Models of Geographical Systems*. Springer.
- Alger, I., and Weibull, J. W. (2013). Homo Moralis – Preference Evolution Under Incomplete Information and Assortative Matching. *Econometrica*. 81 (6), pp.2269-2302.
- Amaral, M. A., and Javarone, M. A. (2018). Heterogeneous Update Mechanisms in Evolutionary Games: Mixing Innovative and Imitative Dynamics. *Physical Review E*. 97 (4), Article 042305. URL: <https://journals.aps.org/pre/abstract/10.1103/PhysRevE.97.042305>.
- Andreoni, J. (1990). Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving. *The Economic Journal*. 100 (401), pp.464-477.
- Andreoni, J., and Croson, R. (2008). Partners Versus Strangers: Random Rematching in Public Goods Experiments. In: Plott, C. R., and Smith, V. L. (Eds.). *Handbook of Experimental Economics Results, Volume 1*. Elsevier.
- Andreoni, J., and Miller, J. H. (1993). Rational Cooperation in the Finitely Repeated Prisoner's Dilemma: Experimental Evidence. *The Economic Journal*. 103 (418), pp.570-585.
- Anderson, P. W., Arrow, K. J., and Pines, D. *The Economy as an Evolving Complex System*. Westview Press.
- Apicella, C. L., and Silk, J. B. (2019). The Evolution of Human Cooperation. *Current Biology*. 29 (11), pp.R447-R450.
- Ariely, D. (2010). *Predictably Irrational*. Harper Collins USA.
- Arthur, W. B. (1994). Inductive Reasoning and Bounded Rationality. *The American Economic Review*. 84 (2), pp.406-411.
- Arthur, W. B. (2013). *Complexity Economics: A Different Framework for Economic Thought* (SFI Working Paper: 2013-04-012). Santa Fe Institute.
- Arthur, W. B. (2021). Foundations of Complexity Economics. *Nature Reviews Physics*. 3 (2), pp.136-145.
- Arthur, W. B., Durlauf, S. N., and Lane, D. A. (1997). *The Economy as an Evolving Complex System II*. Addison-Wesley.
- Arthur, W. B., Holland, J. H., LeBaron, B., Palmer, R. G., and Tayler, P. (1996). *Asset Pricing Under Endogenous Expectations in an Artificial Stock Market* (SFI Working Paper: 1996-12-093). Santa Fe Institute.
- Axelrod, R. (1984). *The Evolution of Cooperation*. BasicBooks.

- Axelrod, R. (1998). *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*. Princeton University Press.
- Axelrod, R. (2006). Advancing the Art of Simulation in the Social Sciences. In: Rennard, J.-P. (Ed.). *Handbook of Research on Nature Inspired Computing for Economics and Management*. Idea Group.
- Axelrod, R., and Hamilton, W. D. (1981). The Evolution of Cooperation. *Science*. 211 (4489), pp.1390-1396.
- Banerjee, A., and Weibull, J. W. (1996). Evolution and Rationality: Some Recent Game-Theoretic Results. In: Allen, B. (Ed.). *Economics in a Changing World: Volume 2 Microeconomics*. Macmillan, pp.90-114.
- Barr, J., and Tassier, T. (2010). Endogenous Neighborhood Selection and the Attainment of Cooperation in a Spatial Prisoner's Dilemma Game. *Computational Economics*. 35 (3), pp.211-234.
- Baumol, W. J., Panzar, J. C., and Willig, R. D. (1988). *Contestable Markets and the Theory of Industry Structure*. Harcourt College Publishers.
- Ben-Ner, A., Putterman, L., Kong, F., and Magan, D. (2004). Reciprocity in a Two-Part Dictator Game. *Journal of Economic Behavior & Organization*. 53 (3), pp.333-352.
- Bernard, A., André, J.-B., and Bredeche, N. (2016). To Cooperate or Not to Cooperate: Why Behavioural Mechanisms Matter. *PLOS Computational Biology*. (May 5, 2016). URL: <https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1004886>.
- Binmore, K., and Shaked, A. (2010). Experimental Economics: Where Next? *Journal of Economic Behavior & Organization*. 73 (1), pp.87-100.
- Bordalo, P., Gennaioli, N., and Shleifer, A. (2017). *Memory, Attention, and Choice* (NBER Working Paper 23256). National Bureau of Economic Research.
- Boyd, R., and Richerson, P. J. (1992). Punishment Allows the Evolution of Cooperation (or Anything Else) in Sizable Groups. *Ethology and Sociobiology*. 13 (3), pp.171-195.
- Boyd, R., and Richerson, P. J. (2009). Culture and the Evolution of Human Cooperation. *Philosophical Transactions of the Royal Society B: Biological Sciences*. 364 (1533), pp.3281-3288.
- Bramoullé, Y., and Kranton, R. (2007). Public Goods in Networks. *Journal of Economic Theory*. 135 (1), pp.478-494.
- Brock, W. A. (2000a). Some Santa Fe Scenery. In: Colander, D. (Ed.) *The Complexity Vision and the Teaching of Economics*. Edward Elgar. pp.29-49.
- Brock, W. A. (2000b). Whither Nonlinear? *Journal of Economic Dynamics and Control*. 24 (5-7), pp.663-678.
- Buchanan, M. (2009). Meltdown Modelling. *Nature*. 460 (7256), pp.680-682.
- Burda, Z., Jukiewicz, J., and Nowak, M. (2003). Is Econophysics a Solid Science? *Acta Physica Polonica B*. 34 (1), pp.87-131.

- Camerer, C., and Ho, T.-H. (1999). Experience-Weighted Attraction Learning in Normal Form Games. *Econometrica*. 67 (4), pp.827-874.
- Cardinot, M., Griffith, J., and O’Riordan, C. (2018). A Further Analysis of the Role of Heterogeneity in Coevolutionary Spatial Games. *Physica A: Statistical Mechanics and its Applications*. 493 (1), pp.116-124.
- Cartwright, E. J., and Lovett, D. (2014). Conditional Cooperation and Marginal per Capita Return in Public Goods Games. *Games*. 5 (4), pp.234-256.
- Charness, G., Gneezy, U., and Kuhn, M. A. (2013). Experimental Methods: Extra-Laboratory Experiments: Extending the Reach of Experimental Economics. *Journal of Economic Behavior & Organization*. 91 (1), pp.93-100.
- Charness, G., and Rabin, M. (2002). Understanding Social Preferences with Simple Tests. *The Quarterly Journal of Economics*. 117 (3), pp.817-869.
- Chen, M., Wang, L., Wang, J., Sun, S., and Xia, C. (2015). Impact of Individual Response Strategy on the Spatial Public Goods Game Within Mobile Agents. *Applied Mathematics and Computation*. 251 (1), pp.192-202.
- Chen, M., Wang, L., Sun, S., Wang, J., and Xia, C. (2016). Evolution of Cooperation in the Spatial Public Goods Game with Adaptive Reputation Assortment. *Physics Letters A*. 380 (1-2), pp.40-47.
- Chick, V. (1998). On Knowing One’s Place: The Role of Formalism in Economics. *The Economic Journal*. 108 (451), pp.1859-1869.
- Cilliers, P. (1998). *Complexity and Postmodernism: Understanding Complex Systems*. Routledge.
- Colander, D. (2000). *The Complexity Vision and the Teaching of Economics*. Edward Elgar.
- Colander, D., Holt, R. P. F., and Rosser Jr., J. B. (2004). *The Changing Face of Economics: Conversations with Cutting Edge Economists*. University of Michigan Press.
- Colander, D., Howitt, P., Kirman, A., Leijonhuvud, A., and Mehrling, P. (2008). Beyond DSGE Models: Toward an Empirically Based Macroeconomics. *The American Economic Review*. 98 (2), pp.236-240.
- Conybeare, J. A. C. (1984). Public Goods, Prisoners’ Dilemmas and the International Political Economy. *International Studies Quarterly*. 28 (1), pp.5-22.
- Corts, K. S. (1998). Third-Degree Price Discrimination in Oligopoly: All-Out Competition and Strategic Commitment. *The RAND Journal of Economics*. 29 (2), pp.306-323.
- Cox, J. C. (2004). How to Identify Trust and Reciprocity. *Games and Economic Behavior*. 46 (2), pp.260-281.
- Dal Bó, P. (2005). Cooperation under the Shadow of the Future: Experimental Evidence from Infinitely Repeated Games. *The American Economic Review*. 95 (5), pp.1591-1604.
- Dal Bó, P., and Frechette, G. R. (2011). The Evolution of Cooperation in Infinitely Repeated Games:

- Experimental Evidence. *The American Economic Review*. 101 (1), pp.411-429.
- Dana, J., Cain, D. M., and Dawes, R. M. (2006). What You Don't Know Won't Hurt Me: Costly (but Quiet) Exit in Dictator Games. *Organizational Behavior and Human Decision Processes*. 100 (2), pp.193-201.
- Dasgupta, P. (2009). Trust and Cooperation Among Economic Agents. *Philosophical Transactions of the Royal Society B: Biological Sciences*. 364 (1533), pp.3301-3309.
- Davidson, P. (1984). Reviving Keynes's Revolution. *Journal of Post Keynesian Economics*. 6 (4), pp.561-575.
- Dawes, R. M., and Thaler, R. H. (1988). Anomalies: Cooperation. *Journal of Economic Perspectives*. 2 (3), pp.187-197.
- Dawkins, R. (2016). *The Selfish Gene* (40th Anniversary Ed.). Oxford University Press.
- De Marti, J., and Zenou, Y. (2015). Network Games Under Incomplete Information. *Journal of Mathematical Economics*. 61 (1), pp.221-240.
- Dong, Y., Zhang, B., and Tao, Y. (2016). The Dynamics of Human Behavior in the Public Goods Game with Institutional Incentives. *Nature: Scientific Reports*. 6, Article 28809. URL: <https://www.nature.com/articles/srep28809>.
- Dow, S. C. (2000). Prospects for the Progress of Heterodox Economics. *Journal of the History of Economic Thought*. 22 (2), pp.157-170.
- Dreber, A., Fudenberg, D., and Rand, D. G. (2014). Who Cooperates in Repeated Games: The Role of Altruism, Inequity Aversion, and Demographics. *Journal of Economic Behavior & Organization*. 98 (February), pp.41-55.
- Enke, B. (2019). Kinship, Cooperation, and the Evolution of Moral Systems. *The Quarterly Journal of Economics*. 134 (2), pp.953-1019.
- Epstein, J. M., and Axtell, R. (1996). *Growing Artificial Societies: Social Science from the Bottom Up*. The Brookings Institution.
- Fehr, E., and Gächter, S. (2000a). Cooperation and Punishment in Public Goods Experiments. *The American Economic Review*. 90 (4), pp.980-994.
- Fehr, E., and Gächter, S. (2000b). Fairness and Retaliation: The Economics of Reciprocity. *Journal of Economic Perspectives*. 14 (3), pp.159-181.
- Fehr, E., and Rockenbach, B. (2003). Detrimental Effects of Sanctions on Human Altruism. *Nature*. 422 (6928), pp.137-140.
- Fehr, E., and Schmidt, K. M. (1999). A Theory of Fairness, Competition, and Cooperation. *The Quarterly Journal of Economics*. 114 (3), pp.817-868.
- Fiala, L., and Suetens, S. (2017). Transparency and Cooperation in Repeated Dilemma Games: A Meta Study. *Experimental Economics*. 20 (4), pp.755-771.

- Fischbacher, U., Gächter, S., and Fehr, E. (2001). Are People Conditionally Cooperative? Evidence from a Public Goods Experiment. *Economics Letters*. 71, pp.397-404.
- Friedman, M. (1953). The Methodology of Positive Economics. In: Friedman, M. *Essays in Positive Economics*. University of Chicago Press. pp.3-46.
- Foster, J. (2006). Why is Economics Not a Complex Systems Science? *Journal of Economic Issues*. 40 (4), pp.1069-1091.
- Fukuyama, F. (1998). Reviewed Work(s): The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration by Robert Axelrod: System Effects: Complexity in Political and Social Life by Robert Jervis. *Foreign Affairs*. 77 (2), pp.142-143.
- Gardner, A., and West, S. A. (2004). Cooperation and Punishment, Especially in Humans. *The American Naturalist*. 164 (6), pp.753-764.
- Gelman, A., and Loken, E. (2013). *The Garden of Forking Paths: Why Multiple Comparisons Can Be a Problem, Even When There Is No "Fishing Expedition" or "P-Hacking" and the Research Hypothesis was Posited Ahead of Time* (Working Paper).
- Gleick, J. (1997). *Chaos: The Amazing Science of the Unpredictable*. Vintage Publishing.
- Gneezy, U., and Rustichini, A. (2000). A Fine is a Price. *Journal of Legal Studies*. 29 (1), pp.1-17.
- Hacking, I. (1967). Slightly More Realistic Personal Probability. *Philosophy of Science*. 34 (4), pp.311-325.
- Hagmann, D., and Tassier, T. (2014). Endogenous Movement and Equilibrium Selection in Spatial Coordination Games. *Computational Economics*. 44 (3), pp.379-395.
- Hamilton, W. D. (1971). Geometry for the Selfish Herd. *Journal of Theoretical Biology*. 31 (2), pp.295-311.
- Harper, D. A., and Lewis, P. (2012). New Perspectives on Emergence in Economics. *Journal of Economic Behavior & Organization*. 82 (2-3), pp.329-337.
- Hatna, E., And Benenson, I. (2012). The Schelling Model of Ethnic Residential Dynamics: Beyond the Integrated - Segregated Dichotomy of Patterns. *Journal of Artificial Societies and Social Simulation*. 15 (1), Article 6. URL: [https://www.jasss.org/15/1/6.html#:~:text=The%20Schelling%20model%20of%20segregation%20\(Schelling%201971%2C%201978\)%20is,148](https://www.jasss.org/15/1/6.html#:~:text=The%20Schelling%20model%20of%20segregation%20(Schelling%201971%2C%201978)%20is,148) .
- Hauert, C., Traulsen, A., Brandt, H., Nowak, M. A., and Sigmund, K. (2007). Via Freedom to Coercion: The Emergence of Costly Punishment. *Science*. 316 (5833), pp.1905-1907.
- He, J., Wang, J., Yu, F., and Zheng, L. (2020). Reputation-Based Strategy Persistence Promotes Cooperation in Spatial Social Dilemma. *Physics Letters A*. 384 (27), Article 126703. URL: <https://www.sciencedirect.com/science/article/pii/S0375960120305703>.
- Heilbroner, R. L. (2000). *The Worldly Philosophers: The Lives, Times, and Ideas of the Great Economic Thinkers* (7th ed.). Penguin Books Ltd.

- Helbing, D. (2009). *Systemic Risks in Society and Economics* (SFI Working Paper: 2009-12-044). Santa Fe Institute.
- Helbing, D., and Balmelli, S. (2011). *How to Do Agent-Based Simulations in the Future: From Modeling Social Mechanisms to Emergent Phenomena and Interactive Systems Design* (SFI Working Paper: 2011-06-024). Santa Fe Institute.
- 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), Article 083005. URL: <https://arxiv.org/abs/1007.0431>.
- Helbing, D., and Yu, W. (2009). The Outbreak of Cooperation Among Success-Driven Individuals Under Noisy Conditions. *Proceedings of the National Academy of Sciences of the United States of America*. 106 (8), pp.3680-3685.
- Helbing, D., and Yu, W. (2010). The Future of Social Experimenting. *Proceedings of the National Academy of Sciences of the United States of America*. 107 (12), pp.5265-5266.
- Heppenstall, A. J., Crooks, A. T., See, L. M., and Batty, M. (2012). *Agent-Based Models of Geographical Systems*. Springer Science.
- Herings, P. J.-J., Peeters, R., Tenev, A. P., and Thuijsman, F. (2021). Naïve Imitation and Partial Cooperation in a Local Public Goods Model. *Journal of Economic Behavior & Organization*. 191 (1), pp.162-185.
- Hirshliefer, J. (1978). Competition, Cooperation, and Conflict in Economics and Biology. *The American Economic Review*. 68 (2), pp.238-243.
- Hobbes, T. (2017 [1651]). *Leviathan*. Penguin Classics.
- Hoffman, E., McCabe, K., and Smith, V. L. (1996). Social Distance and Other-Regarding Behavior in Dictator Games. *The American Economic Review*. 86 (3), pp.653-660.
- Hofweber, T. (2021). Logic and Ontology. In: Zalta, E. N. (Ed.). *The Stanford Encyclopedia of Philosophy* (Spring 2021 Edition). URL: <https://plato.stanford.edu/cgi-bin/encyclopedia/archinfo.cgi?entry=logica-ontology>.
- Holland, J. H. (1992). Complex Adaptive Systems. *Daedalus*. 121 (1), pp.17-30.
- Holland, J. H. (2006). Studying Complex Adaptive Systems. *Journal of Systems Science and Complexity*. 19 (1), pp.1-8.
- Holland, J. H., and Miller, J. H. (1991). Artificial Adaptive Agents in Economic Theory. *The American Economic Review*. 81 (2), pp.365-370.
- Janssen, M. A., and Ahn, T. K. (2006). Learning, Signaling, and Social Preferences in Public-Goods Games. *Ecology and Society*. 11 (2), Article 21. URL: <https://asu.pure.elsevier.com/en/publications/learning-signaling-and-social-preferences-in-public-good-games>.
- Jackson, M. O., Rogers, B., and Zenou, Y. (2016). *The Economic Consequences of Social Network Structure* (IFN Working Paper No. 1116). Research Institute of Industrial Economics.

- Jervis, R. (1988). Realism, Game Theory, and Cooperation. *World Politics*. 40 (3), pp.317-349.
- Judd, K. L. (2005). *Computationally Intensive Analyses in Economics*. Hoover Institution.
- Kahneman, D. (2011). *Thinking, Fast and Slow*. Farrar, Straus & Giroux.
- Kahneman, D., Knetsch, J. L., and Thaler, R. H. (1986). Fairness and the Assumptions of Economics. *The Journal of Business*. 59 (4), pp.S285-S300.
- Kahneman, D., Knetsch, J. L., and Thaler, R. H. (1990). Experimental Tests of the Endowment Effect and the Coase Theorem. *The Journal of Political Economy*. 98 (6), pp.1325-1348.
- Kahneman, D., and Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*. 47 (2), pp.263-292.
- Kandori, M. (2021). *Adjustment Dynamics for Human Players* (Working Paper UTMD-017). The University of Tokyo.
- Kawamura, T, and Tse, T. T. K. (2021). *Intelligence Promotes Cooperation in Long-Term Interaction: Experimental Evidence in Infinitely Repeated Public Goods Games* (Discussion Paper No. 1146). Osaka University: The Institute of Social and Economic Research.
- Keynes, J. M. (1936). *The General Theory of Employment, Interest and Money*. Harcourt.
- Kirchkamp, O., and Nagel, R. (2007). Naive Learning and Cooperation in Network Experiments. *Games and Economic Behavior*. 58 (2), pp.269-292.
- Kreps, D. M., Milgrom, P., Roberts, J., and Wilson, R. (1982). Rational Cooperation in the Finitely Repeated Prisoners' Dilemma Game. *Journal of Economic Theory*. 27 (2), pp.245-252.
- Kube, S., Maréchal, M. A., and Puppe, C. (2012). The Currency of Reciprocity: Gift Exchange in the Workplace. *American Economic Review*. 102 (4), pp.1644-1662.
- Laibson, D. (1997). Golden Eggs and Hyperbolic Discounting. *The Quarterly Journal of Economics*. 112 (2), pp.443-477.
- Lawson, T. (2003). *Reorienting Economics*. Routledge.
- Lawson, T. (2009). The Current Economic Crisis: Its Nature and the Course of Academic Economics. *Cambridge Journal of Economics*. 33 (4), pp.759-777.
- Liu, J., Meng, H., Wang, W., Li, T., and Yu, Y. (2018). Synergy Punishment Promotes Cooperation in Spatial Public Good Game. *Chaos, Solitons & Fractals: Nonlinear Science, and Nonequilibrium and Complex Phenomena*. 109 (1), pp.214-218.
- Lorenz, E. N. (1963). The Predictability of Hydrodynamic Flow. *Transactions of the New York Academy of Sciences*.
- Lucas, P., de Oliveira, A. C. M., and Banuri, S. (2014). The Effects of Group Composition and Social Preference Heterogeneity in a Public Goods Game: An Agent-Based Simulation. *Journal of Artificial Societies and Social Simulation*. 17 (3), Article 5. URL: <https://www.jasss.org/17/3/5.html>.

- Lugovskyy, V., Puzzello, D., Sorensen, A., Walker, J., and Williams, A. (2017). An Experimental Study of Finitely and Infinitely Repeated Linear Public Goods Games. *Games and Economic Behavior*. 102 (1), pp.286-302.
- Machiavelli, N. (2014 [1532]). *The Prince*. Penguin Classics.
- Madani, K. (2010). Game Theory and Water Resources. *Journal of Hydrology*. 381 (3-4), pp.225-238.
- Mandelbrot, B. B. (1963). The Variation of Certain Speculative Prices. *The Journal of Business*. 36 (4), pp.394-419.
- Marks, R. E. (2007). Validating Simulation Models: A General Framework and Four Applied Examples. *Computational Economics*. 30 (3), pp.265-290.
- Masclét, D., Noussair, C., Tucker, S., and Villeval, M. (2003). Monetary and Nonmonetary Punishment in the Voluntary Contributions Mechanism. *The American Economic Review*. 93 (1), pp.366-380.
- Maynard Smith, J. (1964). Group Selection and Kin Selection. *Nature*. 201, pp.1145-1147.
- Mellström, C., and Johannesson, M. (2008). Crowding Out in Blood Donation: Was Titmuss Right? *Journal of the European Economic Association*. 6 (4), pp.845-863.
- Nash, J. (1951). Non-Cooperative Games. *Annals of Mathematics*. 54 (2), pp.286-295.
- Nelson, R. R. (2015). Physics Envy: Get Over It. *Issues in Science and Technology*. 31 (3). URL: <https://issues.org/physics-envy-get-over-it/>.
- Normann, H.-T., and Wallace, B. (2011). *The Impact of the Termination Rule on Cooperation in a Prisoner's Dilemma Experiment* (DICE Discussion Paper, No. 19). Düsseldorf Institute for Competition Economics.
- Nowak, M. A., and May, R. M. (1992). Evolutionary Games and Spatial Chaos. *Nature*. 359 (6398), pp.826-829.
- Noë, R., and Hammerstein, P. (1994). Biological Markets: Supply and Demand Determine the Effect of Partner Choice in Cooperation, Mutualism and Mating. *Behavioral Ecology and Sociobiology*. 35 (1), pp.1-11.
- Poundstone, W. (2021). John von Neumann. *Encyclopedia Britannica*. Last Accessed: 25th January 2022. URL: <https://www.britannica.com/biography/John-von-Neumann>.
- Quan, J., Tang, C., Zhou, Y., Wang, X., and Yang, J.-B. (2020). Reputation Evaluation with Tolerance and Reputation-Dependent Imitation on Cooperation in Spatial Public Goods Game. *Chaos, Solitons & Fractals: Nonlinear Science, and Nonequilibrium and Complex Phenomena*. 131 (1), Article 109517. URL: <https://www.sciencedirect.com/science/article/pii/S0960077919304692>.
- Realpe-Gómez, J., Vilone, D., Andrighetto, G., Nardin, L. G., and Montoya, J. A. (2018). Learning Dynamics and Norm Psychology Supports Human Cooperation in a Large-Scale Prisoner's Dilemma on Networks. *Games*. 9 (4), Article 90. URL: <https://www.mdpi.com/2073-4336/>

9/4/90.

- Richiardi, M., Leombruni, R., Saam, N., and Sonnessa, M. (2006). A Common Protocol for Agent-Based Social Simulation. *Journal of Artificial Societies and Social Simulation*. 9 (1). URL: <https://www.jasss.org/9/1/15.html>.
- Rustagi, D., Engel, S., and Kosfeld, M. (2010). Conditional Cooperation and Costly Monitoring Explain Success in Forest Commons Management. *Science*. 330 (6006), pp.961-965.
- Sachs, J. L., Mueller, U. G., Wilcox, T. P., and Bull, J. J. (2004). The Evolution of Cooperation. *The Quarterly Review of Biology*. 79 (2), pp.135-160.
- Sargent, R. J. (2013). Verification and Validation of Simulation Models. *Journal of Simulation*. 7 (1), pp.12-24.
- Sargent, T. J. (1993). *Bounded Rationality in Macroeconomics*. Clarendon Press.
- Schumpeter, J. A. (1942). *Capitalism, Socialism, and Democracy*. Routledge.
- Schelling, T. C. (1971). Dynamic Models of Segregation. *Journal of Mathematical Sociology*. 1 (2), pp.143-186.
- Schelling, T. C. (2006). *Micromotives and Macrobehavior* (2nd ed.). W. W. Norton & Company, Inc.
- Screpanti, E., and Zamagni, S. (2005). *An Outline to the History of Economic Thought*. Oxford University Press.
- Simmons, J. P., Nelson, L. D., and Simonsohn, U. (2011). False-Positive Psychology: Undisclosed Flexibility in Data Collection and Analysis Allows Presenting Anything as Significant. *Psychological Science*. 22 (11), pp.1359-1366.
- Simon, H. A. (1962). The Architecture of Complexity. *Proceedings of the American Philosophical Society*. 106 (6), pp.467-482.
- Simon, H. A. (1964). *Theories of Bounded Rationality* (CIP Working Paper #66). Carnegie Mellon University.
- Simon, H. A. (1978). Rational Decision-Making in Business Organizations (Nobel Memorial Lecture, 8 December 1978). In: *Economic Sciences*. The Nobel Foundation.
- Smith, A. (2010 [1759]). *The Theory of Moral Sentiments*. Penguin Classics.
- Smith, A. (1982 [1776]). *An Inquiry into the Nature and Causes of the Wealth of Nations*. Penguin Classics.
- Snyder, G. H. (1971). "Prisoner's Dilemma" and "Chicken" Models in International Politics. *International Studies Quarterly*. 15 (1), pp.66-103.
- Szabó, G., and Fáth, G. (2007). Evolutionary Games on Graphs. *Physics Reports*. 446 (4-6), pp.97-216.
- Szolnoki, A., and Perc, M. (2010). Rewards and Cooperation in the Spatial Public Goods Game. *Europhysics Letters*. 92 (3), Article 38003. URL: <https://iopscience.iop.org/article/>


- Tesfatsion, L. (2002). *Agent-Based Computational Economics: Growing Economies from the Bottom Up* (ISU Economics Working Paper No. 1). Iowa State University.
- Tesfatsion, L. (2006). Agent-Based Computational Economics: A Constructive Approach to Economic Theory. In: Tesfatsion, L., and Judd, K. L. (Eds.). *Handbook of Computational Economics, Volume 2*. Elsevier.
- Tesfatsion, L. (2017). *Modeling Economic Systems as Locally-Constructive Sequential Games*. US Department of Energy & the Pacific Northwest National Library.
- Thaler, R. H., and Shefrin, H. M. (1981). An Economic Theory of Self-Control. *The Journal of Political Economy*. 89 (2), pp.392-406.
- Titmuss, R. M. (1970). *The Gift Relationship: From Human Blood to Social Policy*. George Allen and Unwin Ltd.
- Traulsen, A., Semmann, D., Sommerfeld, R. D., Krambeck, H.-J., and Milinski, M. (2010). Human Strategy Updating in Evolutionary Games. *Proceedings of the National Academy of Sciences of the United States of America*. 107 (7), pp.2962-2966.
- Turner, J. R., and Baker, R. (2020). Just Doing the Do: A Case Study Testing Creativity and Innovative Processes as Complex Adaptive Systems. *New Horizons in Adult Education & Human Resource Development*. 32 (2), pp.40-61.
- Von Neumann, J., and Morgenstern, O. (2004 [1944]). *Theory of Games and Economic Behavior* (60th Anniversary Commemorative Edition). Princeton University Press.
- Waltman, L., van Eck, N. J., Dekker, R., and Kaymak, U. (2013). An Evolutionary Model of Price Competition Among Spatially Distributed Firms. *Computational Economics*. 42 (4), pp.373-391.
- Wang, Y., and Chen, T. (2015). Heuristics Guide Cooperative Behaviors in Public Goods Game. *Physica A: Statistical Mechanics and its Applications*. 439 (1), pp.59-65.
- Wang, Y.-L. (2013). Learning Ability Driven by Majority Selection Enhances Spatial Reciprocity in Prisoner's Dilemma Game. *Chaos, Solitons & Fractals: Nonlinear Science, and Nonequilibrium and Complex Phenomena*. 56 (1), pp.96-100.
- Wardil, L., and da Silva, J. K. L. (2013). The Evolution of Cooperation in Mixed Games. *Chaos, Solitons & Fractals: Nonlinear Science, and Nonequilibrium and Complex Phenomena*. 56 (1), pp.160-165.
- Weibull, J. W., Miettinen, T., Kosfeld, M., and Fehr, E. (2020). Revealed Preferences in a Sequential Prisoners' Dilemma: A Horse-Race Between Six Utility Functions. *Journal of Economic Behavior & Organization*. 173 (1), pp.1-25.
- West, S. A., Griffin, A. S., and Gardner, A. (2007). Evolutionary Explanations for Cooperation. *Current Biology*. 17 (16), pp.R661-R672.
- Wheeler, G. (2020). Bounded Rationality. In: Zalta, E. N. (Ed.). *The Stanford Encyclopedia of Philosophy* (Fall 2020 Edition). URL: <https://plato.stanford.edu/entries/bounded-rationality/>.

- Wu, T., Fu, F., Dou, P., and Wang, L. (2014). Social Influence Promotes Cooperation in the Public Goods Game. *Physica A: Statistical Mechanics and its Applications*. 413 (1), pp.86-93.
- Ye, H., Tan, F., Ding, M., Jia, Y., and Chen, Y. (2011). Sympathy and Punishment: Evolution of Cooperation in Public Goods Game. *Journal of Artificial Societies and Social Simulation*. 14 (4), Article 20. URL: <https://www.jasss.org/14/4/20.html>.
- Zelmer, J. (2003). Linear Public Goods Experiments: A Meta-Analysis. *Experimental Economics*. 6 (3), pp.299-310.
- Zhang, Y. (2013). The Impact of Other-Regarding Tendencies on the Spatial Vaccination Game. *Chaos, Solitons & Fractals: Nonlinear Science, and Nonequilibrium and Complex Phenomena*. 56 (1), pp.209-215.
- Zhang, S., Zhang, Z., Wu, Y., Yan, M., and Xie, Y. (2018). Tolerance-Based Punishment and Cooperation in Spatial Public Goods Game. *Chaos, Solitons & Fractals: Nonlinear Science, and Nonequilibrium and Complex Phenomena*. 110 (1), pp.267-272.
- Zhu, C.-J., Sun, S.-W., Wang, L., Ding, S., Wang, J., and Cheng-yi, X. (2014). Promotion of Cooperation due to Diversity of Players in the Spatial Public Goods Game with Increasing Neighborhood Size. *Physica A: Statistical Mechanics and its Applications*. 406 (1), pp.145-154.
- Zhu, H., Ding, H., Zhao, Q.-Y., Xu, Y.-P., Jin, X., and Wang, Z. (2020a). Reputation-Based Adjustment of Fitness Promotes the Cooperation Under Heterogeneous Strategy-Updating Rules. *Physics Letters A*. 384 (34), Article 126882. URL: <https://www.sciencedirect.com/science/article/pii/S0375960120307490>.
- ZHu, H., Guo, H., Zhang, H., Han, Y., Wang, Z., and Chu, C. (2020b). The Role of Punishment in the Spatial Public Goods Game. *Nonlinear Dynamics*. 102 (4), pp.2959-2968.

A Additional Results

A.1 Baseline Simulation Results with Fehr-Schmidt Inequity Aversion Utility Function

Figure 37: *Evolution of the SPGG featuring agents with inequity aversion utility functions.*




The main dilemma in the public goods game is that the monetary payoff from defecting is always higher than the payoff from cooperating, which incentivises agents to defect and not contribute to a greater common good. In the case of inequity aversion utility functions, where agents' welfare is based both on their own monetary payoff and the inequity between them and their neighbours, it is improbable that the inequity aversion utility counteracts the incentive to defect. As Figure 37 shows, the simulation quickly progresses to a state of almost zero cooperation, and the cooperators present are predominantly caused by the random innovation strategy updating. Since the inequity aversion utility function punishes both defectors for being better off than their cooperating neighbours and cooperators for being worse off than their defecting peers, the main difference in utility is the gulf in monetary payoff between the two strategies, which favours defection. Hence, the number of defectors increases, which leads to more defection by a weakening of the 'guilt' felt by defectors, since there is less of a gap between them and their neighbours on average. Thus, the Fehr-Schmidt

inequity aversion utility function is not suitable to simulate real-world cooperation in this setting.

A.2 Baseline Simulation Results with Charness-Rabin Conditional Welfare Utility Function

Figure 38: *Evolution of the SPGG featuring agents with conditional welfare utility functions.*



As in the case of Fehr-Schmidt inequity aversion utilities, Figure 38 shows the game rapidly proceeding towards near-zero cooperation, with the cooperators present likely being mainly due to the innovation strategy updating rule and not imitation. The logic of this utility function is that depending on whether an agent's payoff is higher than their peers or not, they either receive lower utility from their payoff and a bit more from others', or a higher utility from theirs and negative utility from others' payoffs. The problem with basing this conditional welfare purely on monetary payoff is that everyone is almost always better off when they and those around them are defecting, since defecting maximises short-term monetary payoff. The only exceptions are extreme game states of mass cooperation, which are unlikely to be reached due to the defined payoff incentives. Thus, conditional welfare is also not suitable for simulations of cooperation in the context of this thesis.

B MATLAB Program Files

B.1 Scripts

B.1.1 Baseline Simulation Model

The "Baseline_Final.m" MATLAB script is the main file used for creating the baseline simulation of this thesis. It creates the positions of the agents in the lattice and their respective characteristics (utility, strategy, endowment *etc.*) using multiple 3D-arrays where two of the dimensions are used for specifying their location in the lattice and the third is used to store this information over the 1000+ rounds being played. Then, the script utilises a nested 'for loop' to iterate over these 3D-arrays for each round and each agent, playing out the game and then updating utilities and strategies and the conditions for the next round. Also, some key parameters like average proportion of cooperators are saved during the iterations in order to visualize them at the end of the script in multiple charts.

B.1.2 Simulation with Punishment

This script is similar to the baseline, but adding new parameters like punishment fine and cost along with the new strategy class (punishing cooperators) and updating the utility equations for agents accordingly.

B.1.3 Simulation with Reward

This script, like the one with punishment, adds new parameters for costs and rewards along with the new strategy class (rewarding cooperators) and new equations for utility.

B.1.4 Simulation with Reputation

Adding the reputation mechanism means adding new parameters like reputation range and reputation increment to the baseline, but also updating equations to fit with the new mechanism.

B.1.5 Simulation with Policing

As with the other modified versions of the baseline, this script implements new parameters like fines and policing cost and the new strategy class (policing cooperators), while also updating the equations accordingly.

B.1.6 Simulation with Moral Adjustment

The moral adjustment modification of the baseline simulation adds the morality delta parameter, which is then implemented in the nested ‘for loop’ to update agents’ preferences based on their strategy choices.

B.2 Supporting Functions

B.2.1 Neighbours Function

This function is for identifying the 8 agents around an agent in a 2-D square lattice.

B.2.2 Neighbours Max Function

This function is for identifying the maximum level of utility of the 8 neighbours for each agent, for a specific round.

B.2.3 Fermi Equation Function

This function calculates the probability of an agent switching strategy via imitation using the Fermi equation. The input arguments are the utility of the agent from the last round, the utility they would have had if they chose the other strategy and finally the noise parameter, K .