# Optimality and Evolutionary Stability in Social Decision-Making

Anna Sigalou

Submitted in accordance with the requirements for the degree of

*Doctor of Philosophy*

University of Leeds

Department of Statistics

October 2024

# Contents

# List of Figures

# List of Tables

# Acknowledgements

This project started in the middle of lockdown and was conducted between November 2020 and September 2024 -weird circumstances, to say the least. I have so many people to thank for so many different things.

I would like to thank my primary supervisor, Richard Mann, for guiding me through this process and for going at my pace at the (very rough) first year and a half or so of this project, as well as my examiners: Martín López-García and Jamie Wood. I would also like to warmly thank Valentin Lecheval for all his help, as well as Charlie Pilgrim and Andrew Bate. Lastly I want to thank my family for supporting my studies.

I of course also want to thank all of my friends, for being there for me and for making life easier with their company -those in Athens and Stockholm that although far away (and greatly missed) still managed to stay present, and those in Leeds who made a new place feel like home. I have to honourably mention a few: Aris for his immeasurable help navigating academia and life; Rukia who welcomed me in the department and made me feel included; Luis Mario -we started our PhDs together, and I hope to graduate and celebrate them together too; Nils for brightening up the last several months. Last but not least: Chloe, Dimitra, Giorgia P, Georgia M, Villy -because life is more than work and you all make a difference.

The past four years were incredibly weird. A good time overall that I'm glad to be over. Whoever is reading this, I hope you'll enjoy at least parts of the work presented in the following pages.

“I can appreciate the beauty of a flower. At the same time I can see much more about the flower; it’s not just beauty at this dimension of one centimeter, there is also beauty at a smaller dimension, the inner structure.

A science knowledge only adds to the excitement and mystery and awe of a flower”.

- *Richard P. Feynman (abridged)*

# Intellectual Property Statement

The candidate confirms that the work submitted is her own, except where work which has formed part of jointly authored publications has been included. The contribution of the candidate and the other authors to this work has been explicitly indicated below. The candidate confirms that appropriate credit has been given within the thesis where reference has been made to the work of others.

Chapter 5 contains material from the jointly authored publication Sigalou and Mann, 2023:

**Sigalou, A., Mann, R. P.** (2023)
*"Evolutionary stability of social interaction rules in collective decision-making"*, Physical Biology, 20(4), 045003.
This paper was primarily authored by A. Sigalou who was responsible for developing the code, performing the computational analysis and writing. The contributions of R. P. Mann was through advisory and editorial role.

Chapter 6 contains material that was developed by R. P. Mann and J. Voss.

This chapter was primarily developed by A. Sigalou who was responsible for most of the modelling and analysis. The contributions of R. P. Mann and J. Voss were as follows: R. P. Mann contributed to the steady state probabilities of the two-state Markov Chain analysis (subsection 6.3.3). J. Voss contributed the proof for the infinite sequence of same choices (subsection 6.1.2).

# Abstract

Animals need to make important decisions under uncertainty throughout their lives relating to their survival, such as finding resources, avoiding predators or finding safe resting places. Access to social information provides supplementary information to individuals with incomplete personal information and under some conditions can ameliorate their decision-making.

Traditionally, social behaviour is modelled as an observed trait. Here, I assume that social behaviour, i.e. having access to, and utilising social information is an adapted trait: those who are able to make good use of the available information and as a result make more successful decisions, are preferred by evolutionary selection. This piece of work contributes new insight into social behaviour and provides a more neutral context for understanding the occurrence of some commonly observed behaviours.

In chapters 5 and 7 the evolutionary stability of well-mixed groups is explored. I find that sociality evolves in relation to environmental uncertainty and heuristic decision-making rules, while I also establish a necessary constraint on this process. Chapter 6 explores the long-term behaviour of groups employing different decision-making rules, using Markov chains. In chapter 8 I explore the evolution of sociality in relation to the position of an agent in the sequence for unmixed groups, and explore the dynamics between groups with homogeneous behaviour and a single agent with a behaviour different to the collective one. Finally, chapter 9 summarises this work and proposes some directions for future research.

# Chapter 1

# Introduction

Social animals need to make decisions under uncertainty. The underlying principles and mechanisms of these decisions are still largely unknown. This is due mostly to a simple fact: we don't know what those internal processes actually look like.

One thing we do know, is that there are observable advantages to being social; among many other advantages, by being in a group animals have access to indirect information [D. J. T. Sumpter, 2010]. Bees that have spotted a good resource will do a waggle dance to communicate it to the rest of the colony, making it possible for bees that did not spot a resource of this quality to take advantage of it [Biesmeijer and Seeley, 2005]; fish that have failed to spot a predator will learn about it from others in the school that have, thus gaining access to important information regarding their survival [Magurran et al., 1985] in a phenomenon commonly referred to as the many-eyes hypothesis [Pulliam, 1973].

The effective use of this social information has been the focus of many studies, since its effect on decision-making is not straightforward: while it can supplement personal information and help animals make more accurate decisions, it may also mislead them into making worst ones [Giraldeau et al., 2002; Wolf et al., 2013]. Similarly, the way animals appear to use social information is not completely understood; for example several papers agree on the observation that many animal groups tend to become social to a sub-optimal level [Torney et al., 2015; Sigalou and Mann, 2023].

A variety of models have been developed in an effort to understand sociality and the internal mechanisms that make it possible, with many models being heuristic: a rule based on observations is proposed, which is then examined to determine whether it produces realistic group behaviour in silico [Vicsek et al., 1995; Couzin et al., 2002].

Additionally, both approaches define social behaviour as an observed trait; this assumption is necessary in order to understand **what**: what is this behaviour we are noticing across many species. Yet this approach is limited when the question becomes **how**: how do animals become social, and why in the ways and contexts we observe it at?

To address 'how', an approach considering a simple model with few assumptions can be used to explore the evolution of a group and establish what conditions lead to the emergence of different behaviour. Such work has been performed with interesting results, such as in the case of Preston et al., 2010 where they explore the question of how an organism should forage given limited available information. A similar approach is taken in this work to similarly explore how – and under what conditions – sociality evolves.

Social behaviour is hence considered to be an evolved trait rather than a observed one. The focus then shifts from that described previously, to instead understanding how and why social behaviour comes to be. While the tendency to be social is considered a given in this work, animals are not assumed to be social in a specific way a priori; instead, their level of sociality comes to be via their effort to better adapt to their environment. Social agents (representing vertebrates) belong in groups along with other conspecifics. They are able to observe the prior decisions, and have an at least small tendency to be influenced by their actions. Considering different ways of navigating the environment, I examine how their sociality (here, following social information) evolves and if and how it stabilises.

This examination is done through a theoretical, abstract model. Having this approach allows for some new, fairly generalisable conclusions regarding how sociality evolves in respect to environmental uncertainty and two simple but essential evolutionary mechanisms displayed by many vertebrates: adaptation to the environment and resistance to invasion (see Chapter 3 for a more detailed account). It also sheds a different light to some previously observed results (for example, the fact that some social animals tend to become sub-optimal in their use of social information), as now the adaptive framework contextualises them differently, and offers plausible explanations for them. The evolution of social behaviour is addressed in a variety of cases including in relation to environmental uncertainty, decision-making strategy, group size and group type.

The chosen approach is successful in reproducing two commonly observed traits in groups using social information: consensus decisions [Conradt and Roper, 2007] and a tendency to over-sociality [Torney et al., 2015]. In the first part of this project (Chapters 5 and 7) it is made clear that an important factor to the occurrence of both is the employed decision-making strategy and specifically the way that the available social information is processed.

The key difference in the strategies is the effect they have on bias: do they amplify it, and if so how? Social feedback and bias are known to lead to phenomena such as information cascades [Bikhchandani et al., 1992]. This motivates a slightly different question: how do different decision-making strategies affect the long-term behaviour of groups? By modelling different decision-making strategies using stochastic models, it is possible to gain some preliminary results regarding the connection between bias amplification and the occurrence of cascades.

## 1.1 Overview

This thesis is an analysis of the long-term behaviour of groups of social animals that use a probabilistic, Bayesian updating rule to make a decision in a binary setting.

The main work is split into two main questions: how does social behaviour evolve given different circumstances, and how do different decision-making rules lead to different long-term behaviours in groups employing them? The former is addressed through a computational analysis, and focuses on exploring the relationship between the main parameters of the model, as well as the evolutionary dynamics of the groups' sociality. The latter is partly addressed by the computational work, and further analysed mathematically in a separate chapter. The full thesis is structured as follows:

**Chapter 2** provides a brief overview of the existing work on the subject of social decisions in groups of cognitive agents. It covers themes such as rationality, the use of the sequential decision model and social information. Lastly it overviews the main project aims for this work.

**Chapter 3** presents the relevant mathematical concepts and frameworks for this project. These include both frameworks that are used explicitly for later derivations, and the theory underlying some concepts that are used implicitly throughout this project.

**Chapter 4** introduces the setting and the model that is used throughout this thesis. It outlines how the chosen version of the model is derived, followed by a description of the different memoryless decision-making strategies that are considered in this project. Lastly, the concepts of collective measure and sociality are introduced.

**Chapter 5** shows how groups of social agents employing this model evolve their sociality, under the assumption that they adapt to their environment in the long-term. The relationship between sociality and environmental uncertainty is discussed, and a constrained is imposed. Variations of the model are then considered, and it is shown that the existence of an evolutionary stable strategy depends on group size.

**Chapter 6** supplements the computational work of Chapter 5 by providing a mathematical framework to analyse the long-term behaviour of large groups employing versions of the same decision-making model. It uses Markov Chains to demonstrate that under the existence of bias, over time large groups are expected to amplify the original social bias in the system. Furthermore it shows that for large enough bias, social following can be approximately deterministic, with information cascades occurring for one specific variation.

**Chapter 7** expands the analysis of Chapter 5 to the full version of the chosen sequential decision-making setting, where previous decisions are assumed to be perceived as an ordered sequence. Similar to Chapter 5 the evolutionary dynamics of the group are determined, followed by a comparison between the strategy that considers the ordered sequence, and the three memoryless strategies.

**Chapter 8** builds on the work summarised in the aforementioned chapters by analysing the evolution of sociality for individual agents within the group. It shows how this does not coincide with the collective sociality considered thus far, explains how it further depends on the group's composition, and examines the relationship between a single invader and group.

**Chapter 9** summarises the main contribution of the thesis presented in Chapters 5, 7, 6, and 8 and outlines some directions for future work that naturally follow from it.

## Chapter 2

# Existing Work & Project Aims

Over just a century ago, collective displays like bird flocking were attributed to supernatural phenomena such as telepathy [Selous, 1931]. While tales of complex mechanisms still appear in works of fiction, it has been long established that group motion can be explained by considering the emergent behaviour of simple local interactions.

Aoki, 1982 created a simulation that solely considered the actions of approach, avoidance and orientation between the individuals, while a few years later Reynolds, 1987 created the "BOIDS" model, a simulation of a flock of birds inspired by a particle system where birds behaved as individual actors that demonstrate simple behaviours such as collision avoidance, velocity matching and flock centering. These two works showed that simple principles are enough to lead to group behaviour resembling the one observed in nature, without the need for more complex mechanisms, such as central coordination.

The principle of 'simple interactions lead to emergent (complex) collective behaviour' has since become established in fields such as Complex Systems and Collective Behaviour where agents lacking individual complexity engage in repeated interactions within their groups; these equally simple interactions are enough to produce complex patterns at the collective level [Krause et al., 2002; D. J. T. Sumpter, 2010]. A vast body of literature has flourished, spanning themes from collective movement and pattern formation, to animal social networks and collective decisions [D. J. T. Sumpter, 2010].

This thesis belongs broadly in the field of Collective Behaviour, and in particular deals with the matter of collective decisions. This Chapter provides an outline of literature of relevant notions: collective behaviour (migration and decision-making), sequential decisions, cooperation and decision-making strategies.

## 2.1 Collective Motion

### 2.1.1 Models Inspired from Physics

Similarly to the BOIDS model [Reynolds, 1987], many models on collective movement originate from physics. For instance, swarming (which is one of the most studied collective animal processes) is traditionally described using the Vicsek model, originally motivated by the study of active matter in physics [Vicsek et al., 1995].

Foraging is another good example of a collective task modelled as simple particle-like interactions; what appears to be a coordinated task between several individuals of potentially different abilities and needs and without direct channels of communication has been successfully modelled by applying simple interaction rules [Couzin et al., 2005]. The simple agent based model is described by three equations that specify that each individual animal wants to align her direction of travel with that of her conspecifics while attempting to maintain a minimum distance from them while avoiding collision. The model is complete with an additional relation for an informed proportion of individuals.

Additionally to such models being capable of reproducing moving patterns, they have the capacity to also reproduce of adaptation processes. For instance, Wood and Ackland, 2007 created a simple individual-based model for group formation focused on predation and foraging that included phenotypes for the individuals.

### 2.1.2 Beyond Physics: Social Behaviour

Models from physics have been widely used to approach a wide range of collective phenomena –not only collective motion. Conflict resolution [Pinkoviezky et al., 2018], rational decision-making [Galam, 1997], and opinion dynamics and social spreading [Castellano et al., 2009]) have been approached as self-organised systems at criticality [Mora and Bialek, 2011; Munoz, 2018; Gómez-Nava et al., 2023]. This is also the case for the focal point of this thesis: collective decision-making (i.e. the phenomenon where the majority of an originally undecided group of agents commits to a single option), that has been studied as symmetry breaking in physics, a phenomenon often studied with spin models [Romanczuk and Daniels, 2023].

While this body of literature has contributed significant results, thinking of the animals as actors with agency able to process information has been an important addition; being able to distinguish between social and non-social interactions has been shown to be important in understanding interactions during collective movement [Bode et al., 2012]. For instance, Guttal and Couzin, 2010 incorporated the view of individual animals as 'information processing units' to a system of individuals that engage in simple interactions with neighbours, showing that information sharing is essential during collective migration.

It is also important to consider the social network underlying animals groups, since these are also important factors in the groups' movements. Bode et al., 2011 for instance, show how collective motion is also partly due to the preferential interactions within groups, where individuals respond differently to animals they are familiar with.

## 2.2 Collective Decision-making

While cases like bird migration and fish shoaling deal with the problem of collective navigation, there are other cases where the problem is that of collective decision-making. This too is no easy task: the individuals within a group need to decide which activity to perform, when and how [Conradt and Roper, 2003].

Activity synchrony for example requires some individuals to pay a cost to achieve this [Conradt and Roper, 2000]. Since activity synchrony can be costly, this can explain fission-fusion groups [Conradt and Roper, 2000] as well as sex segregation [Conradt, 1998].

### 2.2.1 Sequential Decisions

Sequential decision-making refers to a specific type of collective decision-making where individuals in a group decide one-by-one between two options; each focal individual decides taking into account the previous decisions. Here, instead of considering a multitude of parameters and several neighbours, the focus is on the way a single individual utilises social information. Just like in the case of collective navigation, simple individual interactions are capable of leading to complex collective behaviour. In this setting, the mechanism is as such: the actions of previous individuals leave a stimulus, which the following individuals perceive and use to navigate the environment, overtime leading to a self-organised group performing a collective task [Grassé, 1959].

Take the example of ant colonies: when a foraging ant discovers a food source, she returns to the nest leaving a chemical trail which the other ants can detect and follow to the food source; upon following the trail, they also reinforce it. This reinforcement process amplifies the signal and increases its accuracy [J. L. Deneubourg et al., 1986]. This simple interaction through the pheromone trail is enough for a group of ants to self-organise and successfully navigate between the colony and a food source –with it additionally being able to discover the shortest path between the two [Goss et al., 1989]. So in general, ant colonies create pheromone trails that enable their navigation, with the trail becoming stronger with every new ant taking it. Yet, this process is inherently stochastic: individual ants can make mistakes, or decide to explore instead of exploiting, especially at the early stages of the pheromone trail which leads to trails with more than one paths. Any new ant will choose a path probabilistically, with a probability that is linked to the intensity of the trail in either path [J.-L. Deneubourg et al., 1989].

So the intensity of the pheromone trail signals to the ants how recently and frequently it has been taken in the past. This conveys information about the quality of each of the available options; if one is perceived as ‘best’, then this is the one that will be followed with a higher chance by this decision-maker. Beckers et al., 1993 considered a scenario with two available options, one of which being the best, and phrased the choice of the best one as a quorum response and specifically as a probability dependant on the pheromone intensity of the trail:

$$P(\text{choose left}) = \frac{(k+L)^n}{(k+L)^n + (k+R)^n} \tag{2.1}$$

Here, $L$: number of ants that have taken the left path, $R$: number of ants that have taken the right path, and $n, k$: parameters that control the steepness of the choice function (fitted experimentally in Beckers et al., 1993) and specifically $n$: determines the degree of nonlinearity of the choice, and $k$: the threshold response to the pheromone.

Perna et al., 2012 provided additional insight into the mechanism behind to the response to the pheromone concentration: through a novel – at the time –experimental technique, they determined that the response to the pheromone follows 'Weber's Law', i.e. the relative difference of concentration between the options. Arganda et al., 2012 further proposed a unified framework for this decision-making mechanism which incorporates Weber's Law among others. This proposed model comprises of three parameters: one for the quality of non-social information, one for the quality of social information, and one for the social information itself. Let options $x$ and $y$ where $x$ is best, $n_x$: the number of animals already chosen option $x$, $n_y$: the number of animals already chosen option $y$, $a$: a measure of the non-social information, $s$: a measure of the reliability of the social information and $k$: the relative impact between the information provided for the two options. Then the probability that option $x$ is better that $y$ is defined as such:

$$P(\text{x is good}) = \frac{1}{1 + as^{-(n_x - kn_y)}} \tag{2.2}$$

Building on this, Pérez-Escudero and de Polavieja, 2011 explored the use of social information for animals navigating uncertain environments with binary options using probability matching; in that approach, the probability that individuals choose an option is the probability that this option has of being the best one. They reinforce the existing result that the model fits experimental data [Ward et al., 2008], as well as the claim that this model is capable of incorporating several types of social information. Specifically, they show that it satisfies the case where the decision-maker observes an aggregated number of the previous choices (such as the number on one option minus the number on the other) as well as the case where a dynamic model with dependencies is considered, where the focal decision-maker now observes an ordered sequence. Arganda et al., 2012 further suggest that this decision-making model provides a unifying framework for collective decision-making in animals, as such:

$$P_x = \frac{P(\text{x is good})}{P(\text{x is good}) + P(\text{y is good})} \tag{2.3}$$

### 2.2.2 Uncertainty

J. L. Deneubourg et al., 1986 talked about a reinforced trail's capability of sustaining a self-organised, cohesive ant colony. However this process is not always successful: scouting ants may make mistakes leading the colony into an ant mill [Franks et al., 1991].

The stochastic nature of the binary sequential decision-making model is in fact vulnerable to mistakes given the right conditions. There is always the potential for mistakes as there exists a probability that an animal will choose the least good option which, if followed strongly, can be amplified misleading the group [Pérez-Escudero and de Polavieja, 2011]. This raises an important point regarding the use of the available social information: this information comes in the form of cues, which signal what might be happening. But since animals' sensory abilities are limited those cues can often be misinterpreted by solitary animals leading to false positives and false negatives, potentially leading them to falsely choosing unsafe places or falsely avoiding safe ones [Wolf et al., 2013].

Broadly speaking, there is an established benefit in using social information since personal information is often insufficient or poor. The low quality of personal information can be mitigated by observing others' actions: this provides indirect information about the state of the environment that supplements the personal information and ameliorates the outcomes of decisions [Valone and Templeton, 2002; Dall et al., 2005]. This is demonstrated in a variety of examples throughout taxa, such as in the case of the 'many eyes' effect where the animals benefit from the fact that others around them may spot a predator (who they would otherwise not see themselves) [Pulliam, 1973; Magurran et al., 1985].

But using social information is not straightforward. Not all social information is relevant as it may either be false, or may be originating from agents with different goals [Mann, 2020], while even when social information is relevant, it can be misleading [Giraldeau et al., 2002; Rieucau and Giraldeau, 2009] potentially even leading to poor information cascades [Bikhchandani et al., 1992]. For that reason, the available social information shouldn't always be trusted. So when, and how, is social information to be used?

Intuitively, it might make sense to turn to social information when personal information is insufficient, as making decisions under reliable social influence while having poor personal information has been shown to increase collective accuracy [Jayles et al., 2017]. But as stated above social information isn't always reliable; in fact in many cases the focal agent can assume that others before her are subject to the same uncertainty as her when interpreting ambiguous cues. In settings like the binary decision scenario considered here, whether relying on social information is optimal or not depends on its quality rather than the lack of personal information: when the environmental information is poor, it's better to rely on personal assessments of the available cues rather than on the social information (as this social information is equally unreliable), while when the environmental information is more reliable it's better to rely on social information [King and Cowlishaw, 2007]. This is in line with theoretical models such as the Condorcet's jury theorem [Boland, 1989], according to which groups can make more accurate decisions than solitary decision-makers, provided that they are already accurate enough (specifically, that the individual probability of being accurate is $p > 0.5$).

### 2.2.3 Conflict and Cooperation

Additionally to information uncertainty, animal groups are also subject to conflicts of interest. This can make consensus decisions harder to achieve, leading to the group taking unshared (rather than shared) decisions [Conradt and Roper, 2009]. This research addresses a crucial point: how to approach dynamics within groups? A wide range of behaviours can be observed during interactions, ranging from altruism and cooperation, to selfishness and cheating [West et al., 2007]. *Altruism* refers to a behaviour exhibited by a single individual that infers a cost to her, while conferring a benefit to someone else. *Cooperation* refers to the act of multiple individuals working together and receiving a shared benefit. *Selfishness* refers to a behaviour that is costly to others, but beneficial to the individual that exhibits it. *Cheating* refers to the lack of cooperation from an individual, who nevertheless benefits from others cooperating. Behaviours that benefit others, such as altruism and cooperation have been explained either by showing they can naturally evolve under some circumstances, or by being enforced: this translates to a mechanism that **rewards** cooperation and **punishes** cheating [Trivers, 1971; West et al., 2007].

Since Darwin, 1964 many scientists believed that its principle of ‘survival of the fittest’ is contradictory to the notion of cooperation, leading to the initial dismissal of cooperative phenomena [Axelrod and Hamilton, 1981]. Cooperation was also disregarded as it clashed with the predictions made by rational agent models that predicted ‘selfish’ behaviour [Dawkins, 2016]. However animals behaviour does not align with the predictions made by these models [Kreps et al., 1982]: instead of opting for optimising their own personal circumstances, animals tend to cooperate as documented in a plethora of experimental studies [Andreoni and Miller, 1993; Cooper et al., 1996; Pothos et al., 2011]. To address this gap between theoretical predictions and experimental results, the theories of kin selection [Hamilton, 1964] and reciprocation [Trivers, 1971] were developed as extensions to evolutionary theory.

Kin selection re-frames competition as occurring in the gene level instead of the individual level; that way, individuals cooperate with other animals they share their genes with as this will maximise the gene’s fitness even if this means that their personal one will suffer a blow. Under this theory, an individual will sacrifice herself if this sacrifice allows a relative (who shares her genes) to survive. This altruistic behaviour is favoured by natural selection as long as the cost of it on the altruistic individual is lower than the the benefit that the recipient will receive. While the general idea has been mentioned before (e.g. by Darwin, 1964), it was formalised by Hamilton, 1964 and came to be known as *Hamilton’s rule.*

On the other hand, reciprocity offers an explanation for altruistic behaviour between individuals that are not genetically related. Altruistic behaviour can be still naturally selected, as it is favourable for those performing them in the long-run despite there being a temporary cost in doing so [Trivers, 1971]. The way reciprocity allows for altruism to evolve and be maintained in the long-run, is due to the individuals having memory of previous interactions, and the possibility to meet again; then, being altruistic towards someone at a certain point can mean that the roles can be reversed at a future interaction.

This way, the cost of the first interaction will be balanced out by the benefit of the future one [Axelrod and Hamilton, 1981].

More recently, a different lens of looking at animal behaviour has been suggested: that of adaptation. Adaptation refers to populations that change over-time to better respond to their specific environment. Adaptation (along with a few other key processes) is in the core of species evolution while it additionally provides an explanation for the seemingly irrational behaviour we observe in animals [Fawcett et al., 2014]. Animals navigate complex and changing environment; the adaptation to these environments lead to behaviours, including types of decision-making, that make perfect sense in their specific contexts [Houston and McNamara, 1999; Hutchinson and Gigerenzer, 2005; Todd and Gigerenzer, 2007; McNamara and Houston, 2009; Fawcett et al., 2013]. The criticism to the rationality assumption argues that the observed behaviours are the result of an adaptation process and that they are reasonable responses within the context of the animals' lives [Fawcett et al., 2014].

This approach allows for a deeper understanding of animal behaviour and its origins. Take the example of irregular preferences; rational choice theory holds the axiom that rational agents should have transitive and regular preferences as this will lead to a maximised expected benefit [Von Neumann and Morgenstern, 2007] – an axiom that is violated as demonstrated in a range of studies [Waite, 2001; Bateson et al., 2002; Shafir et al., 2002; Latty and Beekman, 2011]. This behaviour though can be interpreted as ecologically rational in repeated choices, as they result from adaptation to heterogeneous environments [Houston et al., 2007; Trimmer, 2013].

### 2.2.4 Decision-making Strategies

The work of Pérez-Escudero and de Polavieja, 2011 discusses how simpler versions of the decision-making model are capable of providing a very good approximation of a more complete model; in other words, in the context of their model an individual considering the difference in numbers between the two available options give comparably good results to her considering the full ordered sequence of prior decisions.

Exactly what strategies animals use is a question that is both interesting and hard to answer. Many different strategies that make use of this social response function have been observed experimentally, ranging from ones with a linear relationship [Perna et al., 2012] to ones with a non-linear one [D. J. Sumpter and Pratt, 2009]. Given that we don't know what the individuals' beliefs about the world are we can only make assumptions based on the principle that these individuals want to maximise their utility, but re also subjected to noise and context [Mann, 2018], and potentially differing goals or preferences [Mann, 2020].

## 2.3 Project Aims

A substantial body of research has established that relatively simple interactions between individuals can produce cohesive groups able to perform complex tasks like self-organised motion [Vicsek et al., 1995; Couzin et al., 2002], group migration [Guttal and Couzin, 2010], conflict resolution [Couzin et al., 2005], or consensus decisions [D. J. Sumpter and Pratt, 2009]. Substantial effort has been made to identify what these 'rules of interaction' are from two perspectives. In one, models are proposed and demonstrated to exhibit the required collective behaviour [Vicsek et al., 1995; Couzin et al., 2002, 2005; D. J. Sumpter, 2006; Strömbom, 2011; Bialek et al., 2012; Romanczuk and Schimansky-Geier, 2012; Vicsek and Zafeiris, 2012; Huepe et al., 2015; Herbert-Read, 2016; Lecheval et al., 2018; Ling et al., 2019]. This fulfils a necessary but not sufficient condition for identifying the appropriate rules, since other models may also exhibit similar collective behaviour. A second approach is to collect empirical data on animal movements and behaviours and use this to directly infer the form of interactions [Herbert-Read et al., 2011; Katz et al., 2011; Gautrais et al., 2012; Pettit et al., 2013]. Combining these two approaches creates a powerful framework for identifying the rules governing interactions [D. J. Sumpter et al., 2012].

However, even if one can specify precisely what interactions occur between individuals, this leaves an open question: from among the set of plausible interactions, why do animals use these rules and not others? Although simple interactions between individuals can clearly lead to functional group behaviours, less is known about their evolution and stability on an individual level in groups of unrelated individuals who cannot be assumed to behave according to a single collective goal. Instead, such animals should evolve to make decisions that serve their own selfish interests such as acquiring food and safety.

Making informed decisions depends on reliable information about the world. That information comes in the form of cues, which indicate the state of the environment. As individual sensory abilities are limited, social animals can make use of 'social information', i.e. information provided by the actions of their conspecifics, as a source of indirect information about the state of the environment [Valone and Templeton, 2002; Dall et al., 2005]. But not all social information is relevant or accurate [Giraldeau et al., 2002; Rieucau and Giraldeau, 2009], and relying too heavily on imitating others can potentially lead to poor information cascades [Bikhchandani et al., 1992]. Therefore we can expect that natural selection will drive animals to adopt specific weightings of private and social information depending on the environment they inhabit so as to maximise the quality of their decisions.

In the area of collective migration, large-scale evolutionary simulations have explored the evolution of interaction rules within a model based on social 'forces' [Wood and Ackland, 2007; Guttal and Couzin, 2010], with selection on the individual level based on navigational accuracy. This not only demonstrated the evolution of rules sufficient to keep a group of agents together as a single 'flock', but also showed the emergence of distinct strategies within the group, characterised by 'leaders' and 'followers'. Importantly, these strategies emerged as a result of individual adaptation under natural selection, rather than being specified in the model itself. Models of collective movement are complex due to the continuous nature of the observable behaviour (motion) and the iterated interactions between individuals over time. As such, it is difficult to make reasoned *a priori* arguments about how animals ought to interact on the move so as to accomplish individual goals, and even evolutionary models such as that above generally work within a heuristic framework of 'social forces' – assuming that agents exhibit force-like attraction, repulsion and alignment interactions, and allowing the strength of these forces to be determined by evolution.

A more mathematically tractable area of collective behaviour can be considered in the form of simple sequential decision-making between discrete options. Recent research has focused on deriving likely interaction rules in such a scenario by considering the behaviour of rational agents [Pérez-Escudero and de Polavieja, 2011; Arganda et al., 2012; Mann, 2018]. One such model, developed in Pérez-Escudero and de Polavieja, 2011; Arganda et al., 2012 has had a considerable influence on empirical work, being used to interpret the observed collective behaviour of fish [Miller et al., 2013; Pérez-Escudero et al., 2013; Mann et al., 2014; Kadak and Miller, 2020], birds [Aplin et al., 2014; Farine et al., 2014] and even humans [Eguíluz et al., 2015; Pérez-Escudero and de Polavieja, 2017]. However, aspects of this model remain unspecified by theoretical arguments and must in each case be fitted to the data available. Furthermore, various assumptions made in the model development allow for the possibility that these rules may be vulnerable to exploitation by animals employing a different strategy. Establishing whether the strategy derived in this model is stable is crucial as foundation for the interpretation of the empirical studies which assume its use by the animals under study.

In this body of work I take the model of Pérez-Escudero and de Polavieja, 2011 as a starting point for considering collective decision making, based on its widespread use in interpreting empirical data. Initially I assume well-mixed groups and describe the conceptual and mathematical basis of this model, highlighting potential vulnerabilities due to non-rational assumptions. I identify the key parameters of the model that are left unspecified in theory, and show that these obey necessary relationships under the assumption that animals make decisions optimally, thus reducing the number of degrees of freedom in the model. I then specify alternative strategies an animal might employ using the same conceptual framework, and explore the stability of the baseline model to invasion by these alternatives.

By considering non-mixed groups, I further analyse the effect of position within the sequence – or in other words, the amount of available social information as this is necessarily altered with the focal decision-maker's positioning – and the evolution of sociality as a function of this. I explore how the evolution of sociality and evolutionary pressure varies with the employed decision-making strategy, and explore the relationship between individual agent attempting to personally optimise sociality and the effect on the group's collective measure. I find that the qualitative relationship is highly dependent on the employed decision-making strategy.

Lastly, a preliminary analytical work is done using a Markov chain analysis. I discover indicators relating to the occurrence of information cascades and find that this again is highly dependent on the employed decision-making strategy as this affects the way and rate at which bias accumulates in the system.

# Chapter 3

# Mathematical Background

Throughout this research project, I consider a commonly used setting (such as in Ward et al., 2008; Pérez-Escudero and de Polavieja, 2011; Arganda et al., 2012; Mann, 2018), where agents in a group need to decide between two options. One of the options represents the best decision, be it food, safety, or some other valuable resource. The agents of the group choose consecutively, so each focal agent is able to observe two things: the environment and the choices made by the previous agents. The focal agent will then process this available information and use it to make an informed decision about what choice to make. The principle underlying decision-making within the concept of my model is Bayesian Estimating.

As it will be shown later (see Chapter 4), the above individual process leads to certain outcomes in the level of the group. These decisions that originally relate to individual success (for instance, when a food resource is found via this process), by extension also relate to the overall group's success too. By using concepts associated with Evolutionary Game Theory and Stochastic Processes it is possible to investigate the group's dynamics only by knowing the individual decision-making rules and making a few careful assumptions about the group's relation to its environment.

## 3.1 Bayesian Updating

Unlike the most commonly used frequentist statistics that represent frequencies, Bayesian statistics represent beliefs. Initially there is a *prior* belief about a model, which then gets updated by incorporating data which in turn lead to a *posterior* updated belief. Two elementary concepts in Bayesian statistics are *conditional probabilities* and the *Bayes theorem.*

A *conditional probability* $P(A|B)$ denotes the probability of event $A$ occurring **given** that event $B$ has already occurred. We have that:

$$P(A|B) = \frac{P(A \cap B)}{P(B)} \tag{3.1}$$

where $P(A \cap B) = P(B \cap A)$ is the probability that event $A$ **and** event $B$ happen, while $P(B)$ is the probability of event $B$.

Naturally for the conditional probability $P(B|A)$ it is the case that:

$$P(B|A) = \frac{P(B \cap A)}{P(A)} \tag{3.2}$$

From Equations 3.1 and 3.2, it follows that:

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)} \tag{3.3}$$

Equation 3.3 is known as **Bayes rule**. When Bayes rule is used to update a probability based on another event affecting it, the process is called **Bayesian updating**.

## 3.2 Game Theory

Game Theory studies interactions that in the simplest case are pairwise and involve two possible behaviours. The interacting individuals make decisions in a specific situation, and based on the pair of behaviours chosen they each receive a payoff. This interaction can be summarised in a payoff matrix: a 2x2 matrix containing the payoff of each of the possible four directions this pairwise interaction can take. A pairwise interaction with a specified payoff matrix is called a *game*, and the interacting individuals are called **players**; each player has a set of **actions** as well as **preferences** over the set of action profiles [Osborne et al., 2004]. An interaction can be either one-off, meaning that the two players don't interact before or after this particular game, or repeated –in which case they interact again, often with memory from the previous interactions.

Game Theory assumes rational agents: this means that they will make choices that will benefit them personally. Each available choice has an attributed utility, while the individuals have utility functions; with this set-up, the individuals are expected to make the decisions that will increase these utility functions [Von Neumann and Morgenstern, 2007]. Depending on the combination of individual choices, the agents will either be in conflict or will cooperate; different games have different outcomes, and the outcomes heavily depend on the setup. In some games, such as the prisoner's dilemma, the interaction dynamics are competitive and the expected rational outcome is mutual defection; in others (such as the stag-hunt game) the interaction dynamics are cooperative and the best outcome requires mutual cooperation.

### 3.2.1 One-off games

Take the example of the *prisoner's dilemma*: two suspects have committed a crime together, and are being interrogated separately. If they both cooperate and cover for each-other and don't confess they receive a minimal sentence of one year in jail (each), if they both defect and betray each other they both get a larger sentence of two years, and if the first cooperates for the second but the second defects, then the one who cooperated receives a sentence of three years while the other one walks free (i.e. receives a sentence of zero years).

So the two people have the set of actions *Cooperate (C), Defect (D).* They also have preferences depending on what the outcomes will be for each case: for a focal suspect the preferred outcome is for her to defect and the other one to cooperate, followed by both of them cooperating, followed by both defecting, and lastly the least proffered is her cooperating while the other suspect defects.

There are then four possible interactions:

- they both defect,
- they both cooperate,
- one defects after the other has cooperated,
- one cooperates after the other has defected.

So each decisions comes with a benefit and a cost; the final outcome for each agent will depend on both which move she chooses, and what move the other agent chooses. The outcomes are the following: if you are prisoner 1, you may either defect or cooperate. In the first case, if prisoner 2 has also defected you will receive two years in prison, while if she has cooperated you only receive one. In the second case, if prisoner 2 has also cooperated you walk free, but if she has defected you receive three years in prison. This pairwise interaction can be summarised in the following payoff matrix:

| | | Player 2 | |
|---|---|---|---|
| | | $C$ | $D$ |
| Player 1 | $C$ | $(1,1)$ | $(0,3)$ |
| | $D$ | $(3,0)$ | $(2,2)$ |

In this scenario the two rational prisoners are trying to get the smallest sentence possible (as this would maximise their utility). Under this assumption, the rational course of action here is to defect. This results from examining which is the best response to the other prisoner's possible strategies; if the other prisoner cooperates, defecting would mean going free instead of serving one year, while if the other prisoner defects defecting would mean serving two years instead of three.

### 3.2.2 Iterated Games

In the one-off Prisoner's Dilemma, the expected outcome is mutual defection. However, in the iterated version of the game cooperation is possible.

Axelrod and Hamilton, 1981 conducted a computer tournament for the prisoner's dilemma where different strategies competed against each-other; the interactions were iterated. The tournament was between a variety of strategies that allowed for defection and cooperation (like in the one-off prisoner's dilemma), but with the added capacity of remembering previous interactions. Out of all the strategies, the one that prevailed was 'TIT FOR TAT': a strategy with memory of only the previous interaction, that always starts with cooperation and afterwards simply copies the other player's last move.

Now cooperation is not only possible, but also robust and stable. A key concept that enables cooperation to arise from the interactions of decision-makers that seek to optimise their personal gains is that of **reciprocity** –now possible because of the game being iterated and the players having memory.

### 3.2.3 Stability

In the prisoner's dilemma, mutually defection is not only the expected rational outcome of the game, but is also an equilibrium point.

The outcomes of each strategy (defect or cooperate) for a prisoner are the following: if you are prisoner 1, you may either defect of cooperate. In the first case, if prisoner 2 has also defected you will receive two years in prison, while if she has cooperated you only receive one. In the second case, if prisoner 2 has also cooperated you walk free, but if she has defected you receive three years in prison.

So in both cases the best strategy for prisoner 1 is to defect, as regardless of what prisoner 2 is employing it will lead to a better payoff compared to cooperating. The same is the case for prisoner 2, making mutual defection the expected outcome since no player can achieve a better payoff by changing strategy. The case where no player can expect to gain more by employing a different strategy is called a *Nash equilibrium.*

## 3.3 Evolutionary Game Theory

**Evolution** is the change of characteristics of a population over time (measured in generations). It relies on selection and replication and is driven by the adaptation of the population in its environment [Darwin, 1964].

Evolutionary Game Theory involves applying some principles from Game Theory to ecological scenarios (involving animals instead of rational agents). The (theoretical) animals play evolutionary games through the framework of either pairwise interactions, or interactions between one individual and the rest of the population [Smith, 1984]. The individuals have certain behaviours or strategies that have a certain fitness; in this context *behaviour* is defined as possible options for what action one can take, while a *strategy* is a behavioural programme, or in other words series of actions which in the simplest case can simply be one behaviour.

Unlike traditional Game Theory, who deals with the dynamic between two agents, Evolutionary Game Theory however deals with the entire population of interacting animals, and explores the spread of the strategies within this population [Hofbauer and Sigmund, 2003]. Individuals with higher fitness reproduce more than those with lower fitness, leading the next generation to have a higher proportion of individuals with the 'successful' traits. For simplicity, asexual reproduction is often assumed; this will also be the case here. When considering the interactions between different strategies within a population, one important question to address is their evolution as they compete with each-other. Assuming that each strategy has a certain fitness at any moment, this fitness will determine how it will be employed in the future [Taylor and Jonker, 1978].

### 3.3.1 Frequency-dependent Selection

Fitness within a population is not constant, but frequency dependent [Nowak, 2019]. Imagine two strategies $A$ and $B$, with respective frequencies $x_A$ and $x_B$ and composition of the population denoted by vector $\overrightarrow{x} = (x_A, x_B)$, while the fitnesses are denoted by $f_A(\overrightarrow{x})$, $f_B(\overrightarrow{x})$.

Then the average fitness is: $\phi = x_A f_A + x_B f_B$, and the selection dynamics are:

$$\begin{aligned} \dot{x_A} &= x_A[f_A(\overrightarrow{x} - \phi)] \\ \dot{x_B} &= x_B[f_B(\overrightarrow{x} - \phi)] \end{aligned} \tag{3.4}$$

Since $x_A + x_B = 1$ we can write $x_A = x, x_B = 1 - x$ and hence:

$$\dot{x} = x(1-x)[f_A(x) - f_B(x)] \tag{3.5}$$

That way, finding the equilibria for this differential equation is straightforward as, by setting $\dot{x} = 0$ we obtain that:

- $x = 0$
- $x = 1$
- $x \in (0, 1)$ such that $f_A(x) = f_B(x)$

with the stability of these equilibria depending on $f_A(x), f_B(x)$ [Nowak, 2019].

### 3.3.2 Replicator Dynamics

Replicator dynamics describe how the frequencies of strategies in a population change over time based on their relative fitness or payoffs [Hofbauer and Sigmund, 2003].

Consider a population of $n$ strategies, with $x_i$ being the frequency of strategy $i$ and $x_j$ that of strategy $j$; the payoff for strategy $i$ when interacting with strategy $j$ is given by

$a_{ij}$. Let $x_i$ denote the frequency of $i$ – the expected payoff of $i$ is then given by:

$$f_i = \sum_{j=1}^{n} x_j a_{ij} \tag{3.6}$$

and the average payoff by:

$$\phi = \sum_{i=1}^{n} x_i f_i \tag{3.7}$$

By equating the payoff with fitness as above, the selection dynamics are in this case:

$$\dot{x_i} = x_i(f_j - \phi), i = 1, 2, ..., N \tag{3.8}$$

which is frequency dependent [Nowak, 2019].

### 3.3.3 Evolutionary Stable Strategy

An *Evolutionary Stable Strategy* (ESS) follows in parts from the concept of stability in traditional Game Theory, and in parts from the work of MacArthur and Hamilton [Hamilton, 1964]. Originally developed to analyse the evolution of phenotypes through pair-wise contests, it later also provided a framework for analysing the competition between individuals and the whole population [Smith, 1984].

It refers to a strategy that can be employed by a population, which cannot be invaded by a rare mutant adopting a different strategy [Smith and Price, 1973; Smith and Parker, 1976; Parker, 1978]. An ESS is not the only stable equilibrium that can be reached [Taylor and Jonker, 1978], it is however the most relevant to this Thesis.

#### Pairwise Invadability Plots

Pairwise invadability plots are plots analysing the dynamics between a population dominated by one strategy, and a mutant using a different strategy. It considers a range of cases of between the population and the mutant using the two strategies, and plots their fitness in the same plot showing whether (and when) the mutant can successfully invade the population's strategy [Geritz et al., 1998].

Pairwise invadability plots provide a good tool for exploring the invasion dynamics including equilibrium points, however they are subject to certain specific assumptions: reproduction is asexual, the offsprings' strategies are identical to that of the parents, mutations are infrequent so that populations can stabilise before a new mutation occurs and strategies are protected against extinction.

### 3.3.4 Two-player Games

Two-player games in Evolutionary Game Theory follow the traditional Game Theory approach of pairwise interactions: two players, $A$ and $B$ interact and depending on their strategies and of the order of the interaction (whether player 1 or player 2 goes first) they get a payoff as such: if player $A$ meets $A$ they both get $a$, if $B$ meets $B$ they both get $d$, if $A$ meets $B$ then $A$ gets $b$ and $B$ gets $c$; this is described in a payoff matrix as such [Nowak, 2019]:

| | | Player 2 | |
|---|---|---|---|
| | | $A$ | $B$ |
| Player 1 | $A$ | $a$ | $b$ |
| | $B$ | $c$ | $d$ |

Evolutionary Game Theory assumes a population of $A$ and $B$ that randomly interact, and equates fitness with the payoffs they accumulate from these interactions. Denote as $x_A$ the frequency of $A$ and $x_B$ the frequency of $B$; these now determine the probability of interacting with an individual using this strategy. The expected payoffs for each strategy are:

$$\begin{aligned} f_A &= ax_A + bx_B \\ f_B &= cx_A + dx_B \end{aligned} \tag{3.9}$$

We can set $x = x_A$ and $1 - x = x_B$ in order to work with a single variable. The selection dynamics are:

$$\dot{x} = x(1-x)[(a-b-c+d)x + b - d] \tag{3.10}$$

There are five possibilities for the selection dynamics:

- $A$ dominates $B$: this happens for $a > c$ and $b > d$ in other words it pays more to play $A$.
- $B$ dominates $A$: this happens for $a < c$ and $b < d$ in other words it pays more to play $B$.
- $A$ and $B$ are bistable: this happens for $a > c$ and $b < d$ and is the case where each strategy is the best response to itself.
- $A$ and $B$ coexist stably: this happens for $a < c$ and $b > d$ and is the case where each strategy out-competes the other one.
- $A$ and $B$ are neutral: this happens for $a = c$ and $b = d$ and is the case where each strategy has the exact same payoff.

### 3.3.5 Games in Infinite Populations

Evolutionary Game Theory also uses games to describe pairwise interactions but instead of assuming rational agents, it assumes *frequency dependant* organisms that evolve according to their payoff matrices.

A population of players with fixed strategies interact in a game. After each encounter they receive a payoff, which contributes to their fitness. Strategies that do well reproduce faster, and strategies that do not get out-competed.

Imagine two strategies, $A$ and $B$. Denote $x_A$ the frequency of $A$ and $x_B$ the frequency of $B$. Also denote $f_A(\overrightarrow{x})$ the fitness of $A$ and $f_B(\overrightarrow{x})$ the fitness of $B$. The selection dynamics are then described as such:

$$\begin{aligned} \dot{x_A} &= x_A[f_A(\overrightarrow{x}) - \phi] \\ \dot{x_B} &= x_B[f_B(\overrightarrow{x}) - \phi] \end{aligned} \tag{3.11}$$

where $\phi = x_A f_A(\overrightarrow{x}) + x_B f_B(\overrightarrow{x})$ the average fitness.

Since there are only two types of strategies in the population, $x_A + x_B = 1$ and we can write $x_A = x$, $x_B = 1 - x$ in order to work with one variable. In that case the fitness of $A, B$ respectively is denoted as $f_A(x), f_B(x)$ and the selection dynamics reduce to:

$$\dot{x} = x(1 - x)[f_A(x) - f_B(x)] \tag{3.12}$$

The differential Equation 3.12 has several equilibrium points corresponding to frequencies of $A$ for which the dynamics between $A$ and $B$ reach a stable point that remains constant throughout the selection process. These equilibrium points occur at $\dot{x} = 0$; an equilibrium point $\dot{x} = 0$ is not necessarily stable [Nowak, 2019].

#### Evolutionary Stable Strategy

Imagine a population of $A$ players and a single $B$ mutant where the selection dynamics are as in 3.10. In that case, the question of interest is when can strategy $A$ be stable from strategy $B$ – or in other other, when can $A$ not be out-competed by $B$. When this happens, then strategy $A$ is an **Evolutionary Stable Strategy** (ESS).

The above will be the case if the fitness of $A$ is greater than the fitness of $B$. For an infinitesimally small fraction of individuals using strategy $B$, assume that $x_B = \epsilon$ and $x_A = 1 - \epsilon$. Then the fitness of $A$ is greater than that of $B$ if:

$$\begin{aligned} a(1 - \epsilon) + b\epsilon &> c(1 - \epsilon) + de \\ \Rightarrow a &> c \end{aligned} \tag{3.13}$$

In the case where $a = c$, then the fitness of $A$ is greater than that of $B$ when $b > d$. So selection will oppose invasion either when $a > c$, or when $a = c$ but $b > d$.

For the case of more than two strategies interacting, we again imagine pairwise interactions between them at a time. Let $S_i, S_j$ be two strategies and $E(S_i, S_j)$ be the payoff that strategy $S_i$ receives when playing against strategy $S_j$ [Nowak, 2019]. There are five ways these interact:

1. $E(S_k, S_k) > E(S_i, S_k) \forall i \neq k$; in this case $S_k$ is a strict Nash equilibrium.
2. $E(S_k, S_k) \geq E(S_i, S_k) \forall i$; in this case $S_k$ is a Nash equilibrium.
3. $E(S_k, S_k) > E(S_i, S_k)$
   or
   $E(S_k, S_k) = E(S_i, S_k)$ and $E(S_k, S_i) > E(S_i, S_i) \forall i \neq k$; in this case $S_k$ is an ESS.
4. $E(S_k, S_k) > E(S_i, S_k)$
   or
   $E(S_k, S_k) > E(S_i, S_k)$ and $E(S_k, S_i) \geq E(S_i, S_i) \forall i \neq k$; in this case $S_k$ is a weak ESS (i.e. it is stable against invasion).
5. $E(S_k, S_k) = E(S_i, S_k)$ and $E(S_k, S_i) > E(S_i, S_i) \forall i \neq k$; in this case $S_k$ is an unbeatable strategy.

### 3.3.6 Games in Finite Populations

In finite size groups stochasticity has a non-negligible impact on the dynamics. For this reason, unlike infinite sized groups that are modelled using differential equations, stochastic equations are more appropriate to model groups of finite size [Nowak, 2019].

In groups of finite size, what contributes to the outcome of a game is frequency dependence and drift. The intensity of selection is also important as it determines the effect (weak or strong) that the outcome of a game will have on fitness.

In finite games the conditions for evolutionary stability are different compared to those for infinite groups due to the effect of stochasticity that is now present – and important. Consider again the game with the following payoff matrix:

| | | Player 2 | |
|---|---|---|---|
| | | $A$ | $B$ |
| Player 1 | $A$ | $a$ | $b$ |
| | $B$ | $c$ | $d$ |

for a population of size $N$. There are $i$ individuals using strategy $A$ and $N - i$ individuals using strategy $B$. Each individual can interact with $N - 1$ others: for each individual using strategy $A$ there are other $i - 1$ others using the same strategy, and for every individual using strategy $B$ there are $N - i - 1$ others using the same strategy.

So an individual using strategy $A$ interacts with another individual using the same strategy with probability $(i-1)/(N-1)$ and with an individual using strategy $B$ with probability $(N-i)/(N-1)$. Similarly, an individual using strategy $B$ interacts with another individual using the same strategy with probability $(N-i-1)/(N-1)$ and with an individual using strategy $A$ with probability $i/(N-1)$. Hence the payoffs for $A$ and $B$ respectively are:

$$\begin{aligned} F_i &= \frac{a(i-1)+b(N-i)}{N-1} \\ G_i &= \frac{ci+d(N-i-1)}{N-1} \end{aligned} \tag{3.14}$$

The expected payoff will contribute to fitness with intensity $w$, where $0 \leq w \leq 1$; for $w=0$ the payoff will not contribute to fitness and for $w=1$ the payoff will determine fitness completely, while for $w \to 0$ is the case of weak selection. The expected fitness taking into account the intensity $w$ are:

$$\begin{aligned} f_i &= 1-w+wF_i \\ g_i &= 1-w+wG_i \end{aligned} \tag{3.15}$$

Consider a Moran process (a discrete-time stochastic process describing the evolution of two phenotypes in a population of constant size, described in more detail in subsection 3.5.1) between $A$ and $B$ with $i$ individuals using strategy $A$. The transition probability of reaching state $i+1$ in one step after $i$ is:

$$p_{i,i+1} = \frac{if_i}{if_i+(N-i)g_i}\frac{N-i}{N} \tag{3.16}$$

and the transition probability of reaching state $i-1$ in one step after $i$ is:

$$p_{i,i-1} = \frac{(N-i)g_i}{if_i+(N-i)g_i}\frac{i}{N} \tag{3.17}$$

while the transition probability of reaching state $i$ in one step after $i$ is simply $1-p_{i,i-1}-p_{i,i+1}$. In the absorbing states $i=0$ and $i=N$ we have $p_{0,0}=1$ and $p_{N,N}=1$. So any group at state $i=0$ or state $i=N$ will remain there, while a group at any state $0<i<N$ will eventually reach one of the two absorbing states.

So now we ask what are the fixation probabilities of strategies $A$ and $B$. Let the ratio $p_{i,i-1}/p_{i,i+1}=g_i/f_i$. Using equations from the Moran process the fixation probability for $A$ is:

$$\rho_A = \frac{1}{1+\sum_{k=1}^{N-1}\prod_{i=1}^{k}\frac{g_i}{f_i}} \tag{3.18}$$

and the ratio of the fixation probabilities is:

$$\frac{\rho_A}{\rho_B} = \prod_{i=1}^{N-1} \frac{g_i}{f_i} \tag{3.19}$$

At the limit of weak selection $w \to 0$ the fixation probability of $A$ becomes:

$$\rho_A \simeq \frac{1}{N} \frac{1}{a - (\alpha N - \beta \frac{w}{6})} \tag{3.20}$$

where $\alpha = a + 2b - c - 2d$ and $\beta = 2a + b + c - 4d$.

Selection favours the fixation of $A$ for $\rho_A > 1/N$, or equivalently $\alpha N > \beta$. This can be written as:

$$a(N-2) + b(2N-1) > c(N+1) + d(2N-4) \tag{3.21}$$

Since we are in the case of finite size, the population size is important for the selection dynamics. For a group of two agents, we get:

$$N = 2 \Rightarrow b > c \tag{3.22}$$

This is the case where the individual using strategy $A$ has payoff $b$ and the one using strategy $B$ has payoff $c$, hence for $b > c$ the one using strategy $A$ is more likely to fixate. For a larger population size we have:

$$a + 2b > c + 2d \tag{3.23}$$

Consider the game where $A$ and $B$ are both the best replies to themselves, i.e. $a > c$ and $b < d$. At the limit of a large population size, if the frequency of $A$ is high then $A$ has a larger fitness than $B$ (and vice versa). The equilibrium point is when the two strategies have equal fitness, $F_i = G_i$. For large $N$ this results in the equilibrium point:

$$x = \frac{d - b}{a - b - c + d} \tag{3.24}$$

which denotes the unstable equilibrium between $A$ and $B$. This leads to the inequality $x < 1/3$: it occurs in a large population $N$ at the weak selection limit $w \to 0$ and it leads to the probability that a mutant $A$ takes over the population being over $1/N$.

Now in the case that $A$ dominates $B$ we have $a > c$ and $b > d$. In this case the inequality always holds as $x < 0$: then the fixation of $A$ will be favoured by selection. Yet for $b < c$ strategy $B$ may still be favoured in a small population; in this case we define $N_c$: for $N < N_c$ selection favours the dominated strategy $B$, otherwise it favours dominating strategy $A$.

### Evolutionary Stability

For finite populations of size $N$, a strategy $B$ is an ESS if:

- selection opposes strategy $A$ opposing strategy $B$ (a single mutant $A$ will have a lower fitness from the fitness of strategy $B$)
- selection opposes strategy $A$ replacing strategy $B$ ($\rho_A < 1/N \forall w > 0$)

These conditions are respectively:

$$\begin{aligned} b(N-1) &< c + d(N-2) \\ a(N-2) + b(2N-1) &< c(N+1) + d(2N-4) \end{aligned} \tag{3.25}$$

For $N = 2$ both conditions reduce to $b < c$. For larger $N$ they lead to $b < d$ and $x > 1/3$ respectively. For small populations the traditional ESS concept is not necessary or sufficient; for large populations it is necessary but not sufficient.

## 3.4 Stochastic Processes

*Stochastic (or probabilistic) processes* refer to mathematical objects defined on the probabilistic space and evolve over time [Bertsekas and Tsitsiklis, 2008] i.e. processes with inherent uncertainty, such as the outcome of a die roll or the chances of it raining tomorrow. For these, we cannot know for sure what the outcome will be; for example, in the case of a fair die roll we cannot know whether we will roll a 2 or a 5, but we know what the possible outcomes are and can specify or estimate the probability of each outcome occurring (for example, $\frac{1}{6}$ in the case of rolling a fair die).

If a stochastic process also satisfies the Markov property, then it's a **Markov process** [Norris, 1998; Bertsekas and Tsitsiklis, 2008]. **Markov property** means *memorylessness*: the next state of the system only depends on the current state, and any additional information would not affect the predictions made. For example, if you have been gambling for a while and currently have X amount of money, the amount of money you'll be left with after your current bet only depends on how much you have now – not how much you came into the game with, or how much you won or lost in the previous rounds. Both a player that initially had 1000 and a player that initially had 10, will have the same probability of winning or losing if they currently each have 20.

Markov processes are fairly simple, but very versatile models. They can describe cases ranging from gambling and stock market fluctuations, to chemotaxis. They are also very powerful, as they can describe the long-term behaviour of a system. The sections below give an overview of the basic concepts, and of two relevant types of Markov processes for this project: the two-state Markov Chain and the Random Walk.

### 3.4.1 Some Definitions

A **stochastic process** is defined as a collection of random variables $X_t$ indexed by time: $\{X(t) : t \in T\}$. $X_t$ is the random variable representing the value of the process observed at time $t$, and $T$ represents the index set (which can be interpreted as time). The random variables $X_t$ take values in the sample space $\Omega$ while the random process $X(t)$ in the state space $\mathcal{S}$.

Stochastic processes can be defined on discrete or continuous time, depending on the cardinality of $T$; then the have a discrete or continuous state space $S$ respectively. Conventionally, continuous time is indexed by the use of $t$ ($X_t$) while discrete by the use of $n$ ($X_n$) indicating the step. So a collection of discrete random variables can be written as $X_0, X_1, X_2, ...$, while a collection of continuous-time random variables is written as $\{X_t\}_{t \geq 0}$. This thesis only deals with discrete time.

Take the example of a dice roll; here a stochastic process can be the sum of the dice rolls $X(n)$ where $X_n$ is the outcome of one dice roll at time $n$. For a 6-sided fair dice, the sample space for each $X_n$ is $\Omega = \{1, 2, 3, 4, 5, 6\}$ the state space for the process $X(n)$ is $S = \{1, 2, 3, 4, 5, ....\}$.

### 3.4.2 Two-state Markov Chains

Markov chains are a collection of discrete-time stochastic processes that have the memoryless property, i.e. the effect of the past in the future of the process is summarised only by the current state. Formally, a discrete-time stochastic process $X_0, X_1, X_2, ...$ is a Markov chain if $\mathbb{P}(X_{n+1} = s | X_0, X_1, X_2, ..., X_n) = \mathbb{P}(X_{n+1} = s | X_n), \forall t \geq 0, \forall s$.

A two-state Markov chain is a Markov chain with only two states, say $i$ and $j$. In its simplest, easiest form, each state is reachable from itself and from the other state. Take the following example: $i$ is the event of having a sunny day, and $j$ is the event of having rain. It's possible to have a sunny day today and a sunny day tomorrow, or a sunny day today and a rainy day tomorrow – or all other combinations. So we can transition from $i$ to $i$, from $j$ to $j$, from $i$ to $j$ and from $j$ to $i$. This example can be formulated as a Markov chain with the following schematic:

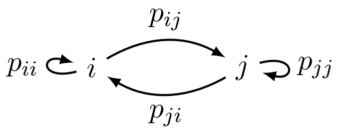

Now it's clear that there are two states, and two transition probabilities from each state. The transition probabilities can be summarised in a *transition probability matrix.* If the transition probabilities $p_{ij}$: *probability of going from state i to state j in one step* and the initial probability vector $X_0$ are known, the Markov chain is fully defined, to the extent that we can use this information to make predictions about the future.

For the example shown above, the transition matrix **A** will be:

$$A = \begin{bmatrix} p_{ii} & p_{ij} \\ p_{ji} & p_{jj} \end{bmatrix} \tag{3.26}$$

Each element of $A$ tells us what is the probability of going from $i$ to $j$ in the next step: so the probability of going from $i$ to $j$ in $n$ steps is the $p_{ij}$th element of $A^n$. So the probability of starting at state $i$, and being in $j$ in 7 steps is given by the element at row $i$ and column $j$ of the matrix $A^7$.

Using the transition matrix, we can also compute the long-term probabilities of finding the system in one of two states, i.e. the probability of being at either of the two states, if we look at the Markov Chain in the very far future (for $n$: large).

### 3.4.3 Random Walks

Consider the integer line $\mathbb{Z}$, and the following process: we're originally at 0, and at each time step $n$ we go up by one (+1) with probability $p$, and down by one (-1) with probability $q = 1 - p$.

This process is known as a *Random Walk* and is formally defined as a stochastic process $X_n$ with discrete time $n = 0, 1, 2, 3, ... = \mathbb{Z}+$ and discrete state space $S = \mathbb{Z}$. At the origin 0: $X_0 = 0$, and then for $n \geq 0$:

$$X_n = \begin{cases} X_n + 1 & \text{with probability } p \\ X_n - 1 & \text{with probability } q \end{cases} \tag{3.27}$$

In its general form, a Random Walk can also be written as follows:

$$X_n = X_0 + \sum_{i=1}^{n} Z_i \tag{3.28}$$

where $X_0 = 0$, and $Z_i$ are independent random variables with $P(Z_i = 1) = p$, $P(Z_i = -1) = q$.

The formalisation of Equation 3.28 is more general, as the state space can be other than $S = \mathbb{Z}$. Furthermore, we can use this structure to calculate the expectation and variance of a Random Walk.

$$\begin{aligned} \mathbb{E}(X_n) &= \mathbb{E}(X_0 + \sum_{i=1}^{n} Z_i) \\ &= \mathbb{E}(X_0) + \sum_{i=1}^{n} \mathbb{E}(Z_i) \\ &= \mathbb{E}(X_0) + n\mathbb{E}(Z_1) \end{aligned} \tag{3.29}$$

For a Random Walk with origin $X_0 = 0 \Rightarrow \mathbb{E}(X_0) = 0$. In that case:

$$\mathbb{E}(Z_1) = \sum_{z \in Z} z\mathbb{P}(Z_1 = z) = 1p + (-1)q = p - q \Rightarrow \mathbb{E}(X_n) = n(p-q) \tag{3.30}$$

$$\begin{aligned} Var(X_n) &= Var(X_0 + \sum_{i=1}^{n} Z_i)) \\ &= Var(X_0) + \sum_{i=1}^{n} Var(Z_i) \\ &= Var(X_0) + nVar(Z_1) \end{aligned} \tag{3.31}$$

For a Random Walk with origin $X_0 = 0 \Rightarrow Var(X_0) = 0$.

### 3.4.4 Exact Distribution of the Random Walk

Let $Y_n$ be the number of upwards steps over the first $n$ steps of the process. Then $Y_n \sim Bin(n,p)$ follows a Binomial distribution with parameters $n, p$.

We know that the Binomial distribution has probability:

$$\mathbb{P}(Y_n = k) = \binom{n}{k} p^k q^{n-k} \tag{3.32}$$

for $k = 0, 1, 2, ..., n$, and $\binom{n}{k}$ the Binomial coefficient.

If $Y_n = k$, it means that the process has taken $k$ steps upwards and $n-k$ steps downwards, so now the process is at position $k-(n-k) = 2k-n$. For example, say we're at $Y_5 = 3$: this denotes that we have taken 3 upwards steps over the first 5 steps. So we have taken $Y_5 = 3$ upwards steps, $n-k = 5-3$ downwards steps, and are at position $2k-n = 2*3-5 = 1$.

Since we know the exact distribution, we know the exact probability of being in position $2k-n$:

$$\mathbb{P}(X_n = 2k-n) = \mathbb{P}(Y_n = k) = \binom{n}{k} p^k q^{n-k} \tag{3.33}$$

which in this example is $\mathbb{P}(Y_5 = 3) = \binom{5}{3} p^3 q^2$.

Note that for a Random Walk starting at $X_0 = 0$, after an odd number of steps we are at an odd-numbered state, while for an even number of steps we are at an even-numbered state.

This means that, if the process returns to the origin 0, this can only happen for an even number of steps. For example, it may follow the path '+1, +1, -1, +1, -1, -1' but not the path '+1, +1, -1, +1, -1'.

### 3.4.5 Recurrence & Transience

A state $i$ is called **recurrent** if the probability of ever returning to it after having already having visited once, is equal to one. Alternatively, if this probability is smaller than 1 then the state is called **transient**.

Let $(X_n)$ be a Markov Chain on the state space $\mathcal{S}$. For $i \in \mathcal{S}$, let $m_i$ be the return probability to state $i$:

$$m_i = \mathbb{P}(X_n = i \text{ for some } n \geq 1 | X_0 = i) \tag{3.34}$$

Then if $m_i = 1$ the state $i$ is recurrent: if $m_i < 1$ state $i$ is transient.

When starting at state $i$, the expected number of times that we will visit $i$ again is given by:

$$\mathbb{E}(\# \text{ visits to i} | X_0 = i) = \sum_{n=0}^{\infty} \mathbb{P}(X_n = i | X_0 = i) = \sum_{n=0}^{\infty} p_{ii}(n) \tag{3.35}$$

where $p_{ii}(n)$ denotes the probability of returning to state $i$ after $n$ steps, after starting at state $i$. For a Markov Chain with transition matrix $\mathbf{P}$:

- If $\sum_{n=0}^{\infty} p_{ii}(n) = \infty$ then state $i$ is recurrent
- If $\sum_{n=0}^{\infty} p_{ii}(n) < \infty$ then state $i$ is transient

Suppose state $i$ is recurrent – i.e. when starting at this state it is the case that $m_i = 1$. After we return to $i$ once it will be as if the chain restarts at this state again, due to the Markov property of memorylessness. Imagine that the process returns to the state $i$ infinite times, with $m_i$ each time. Since the number of returns to $i$ is infinite, the expected number of times that $i$ will be visited again is $\sum_{n=0}^{\infty} p_{ii}(n) = \infty$.

Suppose now that $i$ is transient – i.e. when starting at this state it is the case that $m_i < 1$. If we assume that the process returns to $i$ exactly $r$ times, then each of those $r$ times this has happened with a probability $m_i$. while the first time it does not return this happened with a probability $1 - m_i$. So the probability of returning to state $i$ exactly $r$ times is:

$$\mathbb{P}(\# \text{ returns to } i = r) = {m_i}^r (1 - m_i) \tag{3.36}$$

This follows a geometric distribution $Geom(1 - m_i)$. For $Geom(p)$: a Geometric distribution with parameter $p$ we know that the expectation is $\mathbb{E} = (1 - p)/p$, hence in this case the expectation will be:

$$\begin{aligned} \mathbb{E}(\#\text{returns to } i) = & \\ = \sum_{n=0}^{\infty} p_{ii}(n) & \\ = \frac{1-(1-m_i)}{1-m_i} & \\ = \frac{m_i}{1-m_i} & \end{aligned} \tag{3.37}$$

and given that $m_i < 1$ by definition, the expectation is finite.

### Recurrence & Transience of a Simple Random Walk

Whether a random walk is recurrent or transient depends on whether it's symmetric or biased. A symmetric random walk has $p = q = 1/2$ meaning that $m_i = 1$, and then:

$$\sum_{n=0}^{\infty} p_{00}(n) = \frac{m_i}{1-m_i} = \infty \tag{3.38}$$

meaning that it is recurrent. On the other hand for a symmetric random walk with $p \neq q \neq 1/2$ we have that $m_i < 1$ and hence it is transient.

The above can also be deduced by considering the fact that a simple Random Walk can only return to a starting point for an even number of steps. So for $p_{00}^{n}$: the probability of returning to 0 after having started at 0 after $n$ steps we equivalently have $p_{00}^{2n+1} = 0 \ \forall n$, while $p_{00}^{2n} = \binom{2n}{n} p^n q^n$ (since any given sequence of steps of length $2n$ occurs with probability $p^n q^n$ and happens after the process has gone up by $n$ steps, and down by $n$ steps) [Norris, 1998].

To talk about recurrence and transience we need to find what the sum $\sum_{n=0}^{\infty} p_{00}(n)$ is equal to. By Stirling's formula, $n! \sim \sqrt{2\pi n}(n/e)^n$ as $n \to \infty$. Then we can say [Norris, 1998]:

$$p_{00}^{2n}(n) = \binom{2n}{n} p^n q^n = \frac{(2n!)}{(n!)^2}(pq)^n \sim \frac{(4pq)^n}{\sqrt{2\pi n/2}} \tag{3.39}$$

For a symmetric Random Walk $p = q = 1/2$. In this case $4pq = 1$ and $\sum_{n=0}^{\infty} p_{00}^{2n}(n) = \frac{1}{2\sqrt{2\pi}} \sum_{n=0}^{\infty} \frac{1}{\sqrt{n}} = +\infty$. So for $p = q$ the Random Walk is recurrent. On the other hand, for a biased Random Walk $p \neq q \Rightarrow 4pq < 1$. In that case $\sum_{n=0}^{\infty} p_{00}^{2n}(n) = \frac{1}{2\sqrt{2\pi}} \sum_{t=0}^{\infty} (4pq)^n < +\infty$. So for $p \neq q$ the Random Walk is transient.

## 3.5 Stochastic Models for Finite Groups

Evolutionary Game Theory provides a framework that describes evolutionary processes approaching them as deterministic processes [Hofbauer and Sigmund, 2003]. This is applicable to infinite and finite populations; however, finite groups specifically can also be modelled using stochastic models as stochasticity is an important factor [Nowak, 2019].

### 3.5.1 The Moran Process

Assume a population of constant size $N$ and two types of individuals (type $A$ and type $B$) that reproduce at the same rate: at every step, one individual is randomly chosen for reproduction and one for death ensuring that the population remains of constant size.

The evolution of the two types within the population can be described by a discrete-time stochastic process; since the population size is constant, it is enough to track one of the types (say $A$) since the other one can be easily deduced by deducting the number of individuals of type $A$ from the total population size $N$. Let $X(n)$ define the evolution of the number of type $A$ individuals within the population, and $X_n$ be the number of type $A$ individuals at time $n$. Each time-step $n$ in the parameter space $T$ corresponds to a birth-death event, while the state space $S = (1, 2, ..., N)$ for $N$: the population size. This process is called a **Moran process** [Nowak, 2019].

For a population of size $N$ with $i$ individuals of type $A$, there are $N - i$ individuals of type $B$. In such a case there is a probability of $i/N$ of choosing an individual of type $A$ and a probability of $(N - i)/N$ of choosing an individual of type $B$. According to the reproduction process, there are four possibilities for the next step:

1. An individual of type $A$ is chosen for both reproduction and death. This happens with a probability of $(i/N)^2$, and as a result $X_{n+1} = X_n$ i.e. at the next step is the same as in the current one.
2. An individual of type $B$ is chosen for both reproduction and death. This happens with a probability of $((N - i)/N)^2$, and as a result $X_{n+1} = X_n$.
3. An individual of type $A$ is chosen for reproduction and an individual of type $B$ is chosen for death. This happens with a probability of $(i(N - i)/N^2)$, and as a result the variable $X_{n+1} = i + 1$.
4. An individual of type $B$ is chosen for reproduction and an individual of type $A$ is chosen for death. This happens with a probability of $(i(N - i)/N^2)$, and as a result $X_{n+1} = i - 1$.

In general, by describing the state of the system by the variable $X_n = i$, the probabilities of moving from state $i$ to a new state $X_{n+1} = j$ (where $j$ is either $i + 1$ or $i - 1$) are given by a tridiagonal $(N + 1)$ by $(N + 1)$ transition matrix $P$ with entries:

$$\begin{aligned} p_{i,i-1} &= \frac{i(N-i)}{N^2} \\ p_{i,i} &= 1 - p_{i,i-1} - p_{i,i+1} \\ p_{i,i+1} &= \frac{i(N-i)}{N^2} \end{aligned} \tag{3.40}$$

and all other entries equal to zero. Moreover, assume that once the population has been left with only one type of individual, it has reached an absorbing state:

$$\begin{aligned} p_{0,0} = 1, p_{0,1} = 0 \forall i > 0 \\ p_{N,N} = 1, p_{N,i} = 0 \forall i < N \end{aligned} \tag{3.41}$$

This means that once only individuals of either type $A$ or type $B$ are left, there is no probability that at the next step an individual of the other type will be present. All other states $0 < i < N$ are transient, meaning that the process remains in them for a limited amount of time and over-time eventually reaches one of the two absorbing states of only individuals of type $A$ or only individuals of type $B$ – where it remains forever.

The question to ask then at this point is: for a population starting with $i$ individuals of type $A$, what is the probability that eventually only individuals of type $A$ are left (in other words, that the process has reached the absorbing state $N$)?

To answer this, let $x_i$: the probability of ending up in state $N$ when starting at state $i$. In the case that $i = 0$, the process is already in the absorbing state where only individuals of type $B$ exist, from which the state $N$ is unreachable – hence $x_0 = 0$; on the other hand, in the case that $N$, the process is already in the absorbing state where only individuals of type $A$ exist, which is the state we are interested in reaching – hence $x_N = 1$. For all other states $0 < i < N$, the process can reach state $N$ in three different ways:

- The next step will lead to state $i-1$; from there, the process will reach state $N$ with probability $x_{i-1}$.
- The next step will lead to state $i$; from there, the process will reach state $N$ with probability $x_i$.
- The next step will lead to state $i+1$; from there, the process will reach state $N$ with probability $x_{i+1}$.

So a population of originally $i$ individuals of type $A$, will reach state $i = N$ according to the following recursive process:

$$\begin{aligned} x_0 = 0 \\ x_i = p_{i,i-1}x_{i-1} + p_{i,i}x_i + p_{i,i+1}x_{i+1} \\ x_N = 1 \end{aligned} \tag{3.42}$$

### 3.5.2 General Birth-Death Processes with Neutral Drift

The Moran process is a special case of a **birth-death** process. A birth-death process is a 1-dimensional stochastic process on discrete state space $S = \{0, 1, ..., N\}$. In each event, the random variable $x_i$ can either increase or decrease by 1.

Imagine again a population of fixed size $N$, with $i$ individuals of type $A$ and $N - i$ individuals of type $B$, with equal reproduction rates for $A$ and $B$.

Let $\alpha_i$ be the probability of transitioning from state $i$ to state $i+1$ in one step, and let $\beta_i$ be the probability transitioning from state $i$ to state $i-1$ in one step; note that $\alpha_i + \beta_i \leq 1$ (accounting for the possibility of transition from $i$ to $i$, which occurs with probability $1 - \alpha_i - \beta_i$) [Nowak, 2019].

Consider again a process with two absorbing states, $i = 0$ and $i = N$, and consequently $\alpha_0 = 0$, $\beta_N = 0$. Similarly to the Moran process described above, the probability of ending up at state $N$ when starting at state $i$ is given by the recursive relation:

$$\begin{aligned} x_0 &= 0 \\ x_i &= \beta_i x_{i-1} + (1 - \alpha_i - \beta_i) x_i + \alpha_i x_{i+1} \\ x_N &= 1 \end{aligned} \tag{3.43}$$

Consider now the variable $y_i = x_i - x_{i-1}$, $0 < i < N$ and let $\gamma_i = \frac{\beta_i}{\alpha_i}$. Applying these to the recursion 3.43 this becomes:

$$\begin{aligned} x_i &= \beta_i x_{i-1} + (1 - \alpha_i - \beta_i) x_i + \alpha_i x_{i+1} \\ &\Rightarrow y_{i+1} = \gamma_i y_i \end{aligned} \tag{3.44}$$

And from 3.44 we get:

$$\begin{aligned} y_1 &= x_1 \\ y_2 &= \gamma_1 y_1 = \gamma_1 x_1 \\ y_3 &= \gamma_2 y_2 = \gamma_2 \gamma_1 x_1 \\ &. \\ &. \\ &. \end{aligned} \tag{3.45}$$

where $\sum_{j=1}^{N-1} y_i = 1$. and by summing all of the $y_i$ terms:

$$x_1 = \frac{1}{1 + \sum_{j=1}^{N-1} \prod_{k=1}^{j} \gamma_k} \tag{3.46}$$

and by using $x_i = x_1(1 + \sum_{j=1}^{i-1} \prod_{k=1}^{j} \gamma_k)$:

$$x_i = \frac{1 + \sum_{j=1}^{i-1} \prod_{k=1}^{j} \gamma_k}{1 + \sum_{j=1}^{N-1} \prod_{k=1}^{j} \gamma_k} \tag{3.47}$$

We can now talk about **fixation probabilities**, in other words about the probability that a single type $A$ individual can take over a homogeneous population of type $B$ individuals (or vice versa).

Let $\rho_A$ be the fixation probability of a type $A$ individual in a population of type $B$

individuals, and let $\rho_B$ be the fixation probability of a type $B$ individual in a population of type $A$ individuals. Then $\rho_A = x_1$ and $\rho_B = 1 - x_{N-1}$; hence it is the case that:

$$\rho_A = \frac{1}{1 + \sum_{j=1}^{N-1} \prod_{k=1}^{j} \gamma_k} \tag{3.48}$$

$$\rho_B = \frac{\prod_{k=1}^{N-1} \gamma_k}{1 + \sum_{j=1}^{N-1} \prod_{k=1}^{j} \gamma_k} \tag{3.49}$$

The ratio of the fixation probabilities is $\frac{\rho_B}{\rho_A} = \prod_{k=1}^{N-1} \gamma_k$. If $\frac{\rho_B}{\rho_A} > 1$ it is likely that type $B$ will be fixed over time, while in the case where $\frac{\rho_B}{\rho_A} < 1$, type $A$ will be fixed over time.

### 3.5.3 General Birth-Death Processes with Random Drift

In the previous subsection the assumption was that the two types, $A, B$ reproduced at the same rate – a case described as *neutral drift*. When the two types reproduce at different rates, it is known as a case with **random drift**.

Assume that type $A$ has fitness $r$ and type $B$ has fitness 1. Three cases are possible:

- $r = 1$: selection does not favour one type over the other *neutral drift*
- $r > 1$: selection favours type $A$
- $r < 1$: selection favours type $B$

This fitness difference can be incorporated in the birth-death process by implementing different probabilities of choosing types $A$ and $B$ for reproduction. Let now the probability of choosing an individual of type $A$ for reproduction be $ri/(ri+N-i)$ and the probability of choosing an individual of type $B$ for reproduction be $(N-i)/(ri+N-i)$. The probabilities that they will be chosen for elimination are equal as before, that is $i/N$ for an individual of type $A$ and $(N-i)/N$ for an individual of type $B$.

Now the transition matrix for the case of random drift is:

$$\begin{aligned} p_{i,i-1} &= \frac{N-i}{ri+N-i}\frac{i}{N} \\ p_{i,i} &= 1 - p_{i,i-1} - p_{i,i+1} \\ p_{i,i+1} &= \frac{ri}{ri+N-i}\frac{N-i}{N} \end{aligned} \tag{3.50}$$

The proportion parameter is now:

$$\gamma_i = \frac{p_{i,i-1}}{p_{i,i+1}} = \frac{1}{r} \tag{3.51}$$

and the absorption probability in state $N$ is now given by:

$$x_i = \frac{1 - 1/r^i}{1 - 1/r^N} \tag{3.52}$$

So finally the fixation probabilities are:

$$\begin{aligned} \rho_A = x_1 &= \frac{1 - 1/r}{1 - 1/r^N} \\ \rho_B = 1 - x_N &= \frac{1 - r}{1 - r^N} \end{aligned} \tag{3.53}$$

The ratio of fixation probabilities is $\frac{\rho_B}{\rho_A} = r^{1-N}$ and in the case of an advantageous mutant $A$ $(r > 1)$ and $N >> 1$ this is approximated by $\rho_A = 1 - 1/r$ [Nowak, 2019].

## 3.6 Discussion

This chapter provided a summary of mathematical concepts and frameworks that are relevant to this thesis. It presented and provided definitions and some necessary derivations on stochastic processes, since these are necessary for understanding Chapter 6. It also briefly defined Bayesian updating, which is relevant to the model presented in Chapter 4 and used as a basis for the analyses shown in Chapters 5, 7 and 8. Lastly, it summarised Evolutionary Game Theory for groups of both finite and infinite size; while this framework is not explicitly used in this work, concepts tied to it are referred to throughout (such as the definitions of an evolutionary stable behaviour), while it is argued in Chapter 9 that it can be used to conduct a supplementary analysis as part of future work stemming from this thesis.

# Chapter 4

# Model of Sequential Decision-making

Animals need to decide between distinct options, often in groups. The number of options may vary but in the simplest case there are only two available ones to choose from. This setting has been widely employed in works investigating social decision-making [Pérez-Escudero and de Polavieja, 2011; Arganda et al., 2012; Mann, 2018; Sigalou and Mann, 2023]. While the simplicity of a binary decision is an obvious advantage, recent research implies that it is reasonable as it is proposing that even when presented with multiple options, a multi-choice decision is broken into a series of binary ones [Sridhar et al., 2021].

This chapter provided a methodological introduction, necessary for the following four chapters. It initially describes the model of sequential decision-making as it appears in Pérez-Escudero and de Polavieja, 2011, by repeating the derivation found in that publication for the general case and by showing two of the versions that will also be considered here. It then builds upon this model by introducing two alternative variations for the model (including showing the derivation of one of them), introducing a metric for measuring the collective outcome and by commenting on the concept of sociality – a key concept throughout this project. Note that while Pérez-Escudero and de Polavieja, 2011 denote the available options as $X$ and $Y$, $B$ and $A$ is used here instead as this notation has been preferred in more recent literature, such as in Mann, 2018, 2020; Sigalou and Mann, 2023.

The setting is the following: a group of agents is queuing in a crossroad. The agents sequentially choose between the two available options. The first agent makes a decision independently, having only the environmental information at hand; all following agents additionally have social information, i.e. what the previous agents chose (Figure (4.1)).

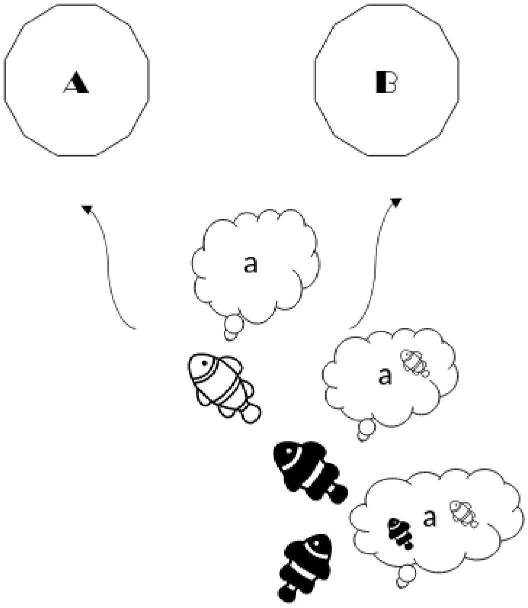

Figure 4.1: Setting: a group of agents performing a sequential decision for a binary choice. The agents have to choose between the options $A$ and $B$ and I assume that $A$ is the best option out of the two. To make her decision, each focal agent observes the environment which has a level of uncertainty (denoted by $a$), as well as the choices of the previous agents.

Each agent decides individually. After all agents in the group have made their respective *individual decisions*, the group has a *collective outcome*. This chapter introduces the individual decision-making model (), and explains how the collective outcome is defined and measured. It also introduces the concept of *decision-making strategies* and *sociality* and their effect in the outcome of decision-making.

## 4.1 Existing Work

This section outlines the individual sequential decision-making process as it appears in Pérez-Escudero and de Polavieja, 2011: it starts by defining the problem, showing the derivation of the generalised version of the rule, and presenting two of the versions that are considered later in the thesis.

### 4.1.1 Individual Decision-making

The problem is as such: assume a setting where the available options are $A$ and $B$, where $A$ is the best option out of the two. The decision is performed according to a probabilistic rule, which in its generalised form is as follows:

$$P_A = \frac{1}{1 + aS} \tag{4.1}$$

where $P_A$ is the probability to choose option $A$ when that is the best option [Pérez-Escudero and de Polavieja, 2011]. In this model $a$ is defined as the non-social parameter as it depends only on the non-social information while $S$ depends on the actions of other individuals; in other words, $a$ describes the uncertainty of the environment and $S$ describes the level of the agents' sociality.

### 4.1.2 Derivation of the Generalised Decision Model

Equation 4.1 is a version of the Bayesian estimation model as developed by Pérez-Escudero and de Polavieja, 2011. The model refers to the above-mentioned setting, and considers each individual making a Bayesian estimation of the probability that each option is the best choice.

Each individual estimates the probability that each choice is the best one based on its non-social information $C$, and the behaviour of the other individuals $I$, in order to decide what behaviour to perform.

Let $A, B$ be the two available options the individuals choose from; then, the probability of $A$ being the best one is $P(A|C, I)$, and the probability that $B$ is the best one is $P(B|C, I) = 1 - P(A|C, I)$. By using Bayes' theorem, $P(A|C, I)$ is given by:

$$P(A|C, I) = \frac{P(I|A, C)P(A|C)}{P(I|B, C)P(B|C) + P(I|A, C)P(A|C)}, \tag{4.2}$$

which can be expressed in the following simplified form:

$$P(A|C, I) = \frac{1}{1 + aS} \tag{4.3}$$

where:

$$a = \frac{P(B|C)}{P(A|C)} \tag{4.4}$$

and

$$S = \frac{P(I|B, C)}{P(I|A, C)} \tag{4.5}$$

In this model $a$ is defined as the non-social parameter as it depends only on the non-social information $C$ while $S$ depends on the actions of other individuals; in other words, $a$ describes the uncertainty of the environment and $S$ describes the level of the agents' sociality.

By definition both parameters are constrained to be $a, S > 0$. Increased values of $a$ correspond to increase in uncertainty as this means that $P(B|C)$ is significantly larger than $P(A|C)$ according to Equation 4.4, while increase in the value of $S$ corresponds to increase in social behaviour. More specifically, $a = 1$ is the case where none of the two available options are favoured over the other, $a < 1$ is the case where option B is favoured over A when A is the best, and $a > 1$ is the case where option A is favoured over B when A is the best. Additionally, $S = 1$ is the case where the focal agent ignores the available social information, $S < 1$ is the case where the focal agent avoids the available social information, and $S > 1$ is the case where the focal agent follows the available social information.

Since in this thesis I am considering the social behaviour of decision-making agents, I will be considering values of $a \in (0, 1)$ as this corresponds to cases where $P(B|C) \leq P(A|C)$ (i.e. there is some bias in favour of option $B$, despite option $A$ being the best one), and $S > 1$ as this corresponds to the case where the focal agent actively follows the available social information.

The value of $S$ depends on the amount of social information available to the decision-maker, which largely depends on how many agents have chosen before the focal; for a fixed value of $a$, the value of $P_A$ will be different for each agent within the group depending on where they are in the sequence. The range of $S$ and the meaning of its values depends on the decision-making strategy, as $S$ takes different forms in each one.

### 4.1.3 Decision-making Strategies

A *decision-making strategy (or rule)* refers to the way the available social information is processed by the agents. The different rules mainly differ in the amount of information available: here I consider four different rules: the **dependencies** rules, the **aggregated** rule, the **majority** rule, and the **dynamic** rule. In terms of the decision model in Equation 4.1, different decision-making is translated in different definitions and forms of the parameter $S$.

#### Dependencies Rule

The dependencies rule constitutes the fuller version of the Bayesian decision-making rule in the sequential decision-making setting described in Pérez-Escudero and de Polavieja, 2011 in their supplementary appendices; it describes the case, where the focal agent is able to observe the full sequence of ordered decisions before her.

All agents have the same $a$, but the focal agent assumes that all other agents apart from her share a different value $\tilde{a}$. The decision-making rule for the focal agent is:

$$P_A = \frac{1}{1 + aS}, \tag{4.6}$$

where $S$ is defined as the product of all previous decisions, as calculated by the focal assuming that others have $\tilde{a}$ instead of $a$:

$$S = \prod \frac{P_{b_i,B}}{P_{b_i,A}} \tag{4.7}$$

where $P_{b_i,B}$ is the probability that agent $i$ chose option $B$ when $B$ was the best option, and $P_{b_i,A}$ is the probability that agent $i$ chose option $B$ when $A$ was the best option.

The focal agent's assumption of a different $\tilde{a}$ means that she assumes others have better or worst information than her, leading her to relying on social information more or less respectively. As this is an assumption about others' personal information $a$, the range for $\tilde{a}$ is $\tilde{a}$ ∈(0,1). A value of $\tilde{a} < a$ means that the focal agent assumes that others have better information than her, while $\tilde{a} > a$ means that she assumes others have worst information.

Observing the full sequence of decisions also means that now the decisions themselves, as well as how old or recent they are, affect their influence on the focal agent's decision-making: more recent decisions will bare more weight compared to older ones, as they contain more recent information about the environment while old ones may be outdated.

**Aggregate Rule**

As demonstrated above Pérez-Escudero and de Polavieja, 2011 showed that in general $S$ may depend on the full ordered sequence of previous decisions as described in section 4.1.3. However, in the focal analysis a version of the model is presented which makes the simplifying assumption that all decisions prior to the focal agent are independent. In this case, $P(I|A,C)/P(I|B,C)$ reduces to:

$$S = \frac{P(I|B,C)}{P(I|A,C)} = s^{n_A - n_B}, \tag{4.8}$$

where $s$ is a parameter that indicates the relative probability that each agent chooses correctly (i.e. $s = 2$ means that the focal agent assumes each previous decision was twice as likely to be correct as to be wrong), $n_A$ is the number of agents that have chosen A and $n_B$ the number of agents that have chosen B. This independence assumption significantly simplifies the form of the calculation, but at the cost of introducing a false belief to the focal agent, which necessarily compromises the optimality of the subsequent decision.

## 4.2 Novel Work

Building on the work of Pérez-Escudero and de Polavieja, 2011, this section presents the two additional versions of the model that are considered in this thesis (and shows the derivation for one of them, since the other one is trivial). Then, it introduces the collective measure that is being introduced and discusses its properties.

### 4.2.1 Majority Rule

In this case the focal agent does not observe the aggregate number of previous decisions in favour of $A$ and $B$, but instead only observes (or responds to) the direction of the majority decision. In this case the appropriate decision rule is:

$$P_A = \frac{1}{1 + as^{-\mathrm{sign}(\Delta n)}} \tag{4.9}$$

I term this the *majority* decision rule, since the social information is a lower-dimensional simplification of that in Chapter 5.

### 4.2.2 Dynamic Rule

The final variation of model 4.1 I consider is the one where an agent observes only the most recent decision before its own. Here the decision rule is given as:

$$P_A = \frac{1}{1 + as^{-d}} \tag{4.10}$$

where $d = 1$ if the most recent choice was $A$ and $d = -1$ if it was $B$. I term this the *dynamic* decision rule in line with similar usage by Mann et al., 2014 which investigated an analogous model empirically in humbug damselfish.

This rule is motivated by the theoretical findings in the supplementary information of Pérez-Escudero and de Polavieja, 2011 and in Mann, 2018 that more recent decisions should be weighted more strongly by an agent able to fully account for the correlations in previous agents' choices, and by the empirical results of Mann et al., 2014 and Kadak and Miller, 2020 which point to both humbug damselfish and zebrafish responding primarily to the most recent choices of conspecifics.

### 4.2.3 Derivation of the Dynamic Decision Rule

Earlier I showed how the generalised decision-making model of Equation 4.1 is derived by Pérez-Escudero and de Polavieja, 2011, and how the aggregate and majority rules follow from it by consider simplifying assumptions. While the dynamic rule is also considered as another simplified version of the same generalised model, its main assumption (of only paying attention to the most recent decision) sets it apart from the other two simplifications. For this reason, in this part the dynamic rule is derived in a way similar to that of the generalised model.

Consider a system where agents choose between two choices, $A$ and $B$, where each individual estimates the probability that each choice is the best one by considering the non-social information $C$ and the behaviour of other individuals $I$. If we take option $A$ to be the last choice, then the probability of choosing $A$ based on the aforementioned information is:

$$P(A|C, I) \tag{4.11}$$

and since there are only two options, it follows that $P(B|C, I) = 1 - P(A|C, I)$. By using Bayes' theorem, we can write Equation 4.11 as:

$$P(A|C, I) = \frac{P(I|A, C)P(A|C)}{P(I|B, C)P(B|C) + P(I|A, C)P(A|C)} \tag{4.12}$$

Now let:

$$a = \frac{P(B|C)}{P(A|C)}, \tag{4.13}$$

$$S = \frac{P(I|B, C)}{P(I|A, C)}. \tag{4.14}$$

By using Equations 4.13 & 4.14, and $P(B|C, I) = 1 - P(A|I, C)$, we can express $P(A|I, C)$ in the simplified form:

$$P(A|C, I) = \frac{1}{1 + aS} \tag{4.15}$$

By assuming that each behaviour is independent, and that the probability of a set of behaviours is the product of the probabilities of the individual behaviours, we can write that probability as such:

$$P(I|A, C) = \prod_{j=1}^{N} P(i_j|A, C) \tag{4.16}$$

And we can now rewrite Equation 4.14 as:

$$S = \prod_{j=1}^{N} \frac{P(i_j|B, C)}{P(i_j|A, C)} \tag{4.17}$$

For $b_k$: different behaviour classes, we assume that there are no differences between agents performing the same behaviour, i.e. they all share the same $P(b_k|B, C)$ and $P(b_k|A, C)$. Now, for each behaviour $b_k$ we can set

$$s_k = \frac{P(b_k|B, C)}{P(b_k|A, C)} \tag{4.18}$$

and for $n_k$ individuals performing $i_k$ behaviour we have

$$\prod_{k=1}^{n_k} \frac{P(i_k|B, C)}{P(i_k|A, C)} = (\frac{P(i_k|B, C)}{P(i_k|A, C)})^{n_k} \tag{4.19}$$

Now we can write Equation 4.14 as:

$$S = \prod_{j=1}^{L} {s_k}^{n_k} \tag{4.20}$$

where $L$ are the different behaviour classes.

Finally, Equation 4.15 becomes

$$P(A|I,C) = (1 + a\prod_{j=1}^{L} {s_k}^{n_k})^{-1} \tag{4.21}$$

Assume now that there are two behaviour classes, 'choosing A when the previous agent chose A' and 'choosing A when the previous agent chose B'. Every focal agent is considered to be the only one performing this behaviour class at each time, and the two behaviours are mutually exclusive. So, for the two agents being considered at the time of the focal agent's decision (i.e. the focal herself, and the previous agent), $n_1$ and $n_2$ will be either $n_1 = 1$ & $n_2 = 0$, or $n_1 = 0$ & $n_2 = 1$, depending on the choice of the previous agent.

We now consider identical sensory information $P(i_1|B,C) = P(i_1|A,C)$, which in turn leads to $s_1 = {s_2}^{-1} = s$, and the decision making rule for the focal agent in this scenario can finally be expressed as:

$$P(A|I,C) = (1 + a{s_1}^{n_1}{s_2}^{n_2})^{-1} = \begin{cases} (1 + as^{+1})^{-1} & \text{if previous chose A} \\ (1 + as^{-1})^{-1} & \text{if previous chose B} \end{cases} \tag{4.22}$$

### 4.2.4 Collective Outcomes

#### Calculating Individual Outcomes

To obtain the precise probability that an agent in a specific part in the sequence will choose option $A$, we need to know the history of the process thus far as this probability is conditional (so that we can know what $\Delta n$ is, as well as what the probability is that these prior decisions were made instead of alternative ones). For example, imagine a process with $a = 0.3$ and $s = 2$ and a group of $N = 3$ agents; the probability that the third agent chooses $A$ is different depending on whether the two prior decisions where $AA$, $AB$, $BB$ or $BA$. Say that $BB$ is observed; then $\Delta n = n_A - n_B = 0 - 2 = -2$ and hence according to Equation 4.15 the probability of choosing $A$ from this point is $P(\text{Choose A}) = \frac{1}{1+as^{-\Delta n}} = \frac{1}{1+0.3*2^{-(-2)}} = 0.930$. Note however that this is not the probability of her being in $A$, as we also need to take into account the probability of being at that point, i.e. the probability that the two prior decisions have been $BB$. The first decision is made without social information (i.e. $\Delta n = 0$) and the probability of choosing $B$ for that agent is $\frac{1}{1+0.3} = 0.769$; the second agent chooses $B$ after the first agent has chosen $B$ with a probability of $1 - P(\text{Choose A})$ and since the prior decision was $B$ and the second agent observes $\Delta n = n_A - n_B = 0 - 1 = -1$ this probability is $1 - \frac{1}{1+0.3*2} = 1 - 0.625 = 0.375$. So overall the second agent will choose option $B$ after the first chose $B$ with a total probability of $0.769 * 0.375 = 0.288$, and finally the third agent will choose option $A$ after observing $BB$ with a final probability of $0.288 * 0.930 = 0.267$.

Apart from obtaining the above probability, we can also obtain the average probability of an agent choosing option $A$ for a given place in the sequence. For that, we simply need to consider all the possible paths –for instance, for the third agent in the above group these will be $AAA$, $ABA$, $ABB$, $BAA$, $BAB$, $BBA$ and $BBB$– and sum the paths that lead to $A$ for her (so $AAA$, $ABA$, $BAA$, $BBA$) since each sequence occurring is independent of the others meaning that the Total Probability Law can be applied.

### Metric of Collective Outcome

When Equation 5.1 is applied sequentially on all group members on a group with $N$ agents, the group will be divided between the two options and will be in one of $N + 1$ possible configurations, corresponding to the number of agents that have chosen $A$, ranging from 0 to $N$; the probability of each possible configuration depends on the values of $a$ and $s$. The reasoning behind the collective measure is very similar to calculating the average probability an agent in place $i$ will choose $A$. By summing over the final configurations and their probabilities for a specific set of $a$ and $s$, I construct a measure for the group's collective behaviour, defined as $\mathbb{E}(n_A)$, that shows the average number of agents that chose option $A$:

$$\mathbb{E}(n_A) = \sum_{i=1}^{N} i p_i \tag{4.23}$$

where $i$ is the number of agents on option $A$ and $p_i$ the probability of i agents being on option $A$, for agents $i = 1, 2, ..., N$ in the group.

Imagine for example that we want to know the value of $\mathbb{E}(n_A)$ for a group of $N = 3$ with $a = 0.3, s = 2$. We can calculate all possible trajectories for that group by first calculating the probabilities that agent $i = 0$ will choose either option, then the probabilities of agent $i = 1$ choosing either option for the two different choices of agent $i = 0$ e.t.c. When every agent has made her decision, we will have $N + 1$ different possible outcomes (one for each case where $n = 0, 1, .., N$ agents have chosen $A$) whose probabilities we will know, and by summing the products of each probabilities with the number of agents corresponding to each we can obtain $\mathbb{E}(n_A)$.

In the above realisation of the model, the parameters are strictly defined and tied to the system at hand. While this approach was successful in replicating experimental results, it isn't able in its current definition to explain the mechanisms of collective decision-making.

## 4.3 Sociality

Social behaviour is a major focal point in this work. The setting being studied is that of a group of social agents navigating an uncertain environment, and within this context the term 'social behaviour' refers to the use of social information when making a decision, as a mitigating mechanism to the uncertain environmental information.

The group of agents is assumed to posses a key feature of animal groups: adaptation. Adaptation means adopting a behaviour that will make their outcomes as best as possible. In the context of this model (Equation 4.1), that will translate to adapting parameter $s$, as parameter $a$ is a fixed property of the environment – and so is $\Delta n$ and $d$ (depending on the strategy) as the focal agent has no control on what has happened previously.

Social behaviour can help with decision-making; especially in cases where personal information $a$ is reliable, paying more attention to what others have done can increase the probability of making the best choice. Yet as all decisions are probabilistic and prone to mistakes, having a high $s$ value is not optimal. Taken to the infinite limit social behaviour leads to deterministic following, in similar ways in the three aforementioned strategies.

Let $s^{\Delta}$, where $s$ is the social behaviour and $\Delta$ the social information; if $\Delta > 0$:

$$\lim_{s\to\infty} \frac{1}{1+as^{\Delta}} = \frac{1}{\infty} = 0 \tag{4.24}$$

and if $\Delta < 0$:

$$\lim_{s\to\infty} \frac{1}{1+as^{\Delta}} = \frac{1}{1} = 1 \tag{4.25}$$

This means that the more social an agent is, the less attention is paid to personal/environmental information; instead, other agents' choices are copied. When $s \to \infty$, the value of $a$ no longer matters; the only thing that matters is whether the previous choices have biased one side over the other – a trend that is more common for larger values of $a$, but still possible for lower ones. So having an unbounded $s$ is problematic as other agents may make poor choices. We will see later that indeed, both the optimal and evolutionary stable social behaviours are constrained to allow for personal information to remain relevant. The necessary constraints on $s$ are discussed in detail in Chapter 5.

For constrained values of $s$ on the other hand, the value of $a$ is impactful on $P_A$. Parameter $a$, while fixed, can take several values, but as it describes the bias towards the one option, it is always $a > 0$. Moreover, values of $a > 1$ will not be considered here, as a value between $0 < a < 1$ are the ones that make sense given the context and its definition.

The effect of $a$ and $s$ is more widely demonstrated in Figure 4.2. This shows the value of $P_A$ for different combinations of these parameters; three distinct cases are visible:

1. $s < 1$: the agents actively avoid social information
2. $s > 1$ & $a < 1$: the agents follow reliable information

3. $s > 1$ & $a > 1$: the agents follow unreliable information

In case (1) the agents actively ignore the available social information. Case (2) is the case where they follow reliable information – here, we see the observation mentioned previously: up until a point, being more social in the presence of that information increases the probability of making the best decision, but that after a point an increase in social behaviour decreases that probability. Case (3) is the case where unreliable social information is being followed – in that case, in principle, the stronger the information is followed, the lower the probability of making the best decision – but like in case (2), there exists a point when being over-social has the opposite effect.

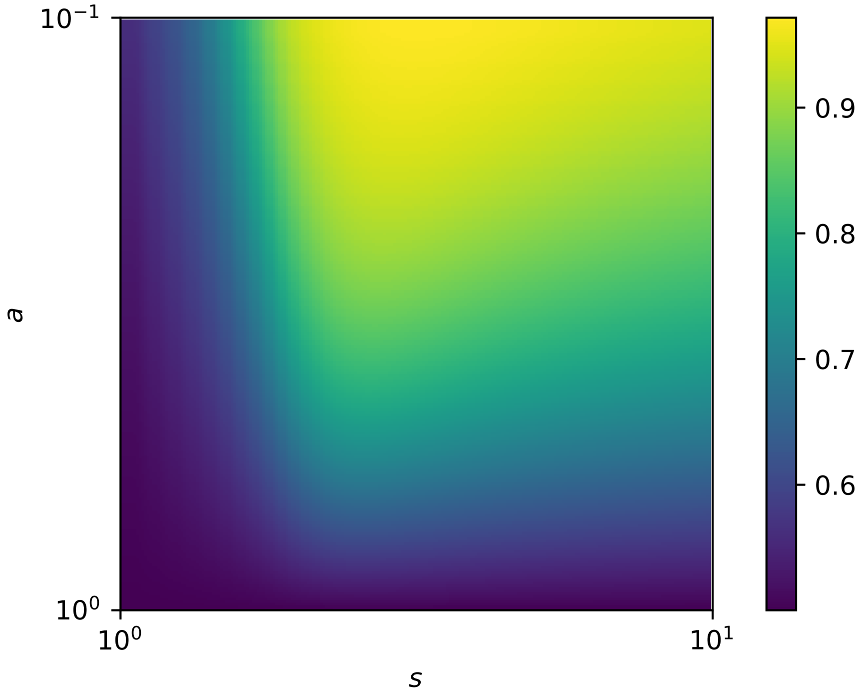

Figure 4.2: Plot showing the value of $P_A$ for different combinations of the parameters $a$ and $s$ for the case of the last agent of a group of $N = 8$ agents employing the aggregate strategy.

## 4.4 Discussion

Systems such as the one studied here are usually modelled as probabilistic; this stochasticity may be attributed either to the agents' limited perception abilities and the ambiguity of the environmental cues [Pratt et al., 2002; Couzin et al., 2005; Guttal and Couzin, 2010], or to the observer's limited perception [Mann, 2018]. Both of these approaches serve the same goal: to introduce noise into the system.

It's irrefutable that noise plays an important role in them; it's exactly that ambiguity of personal information that leads individual agents to source social information from their peers [Dussutour et al., 2009; Faria et al., 2009; Torney et al., 2015]. Acknowledging this feature, the above model is also inherently stochastic. This is achieved due to the probabilistic definition of decision-making, and to the existence of parameter $a$.

Previous work on this setting has worked with specific values for the parameters $a$ and $S$, based on experimental data [Pérez-Escudero and de Polavieja, 2011]. While they might describe what specific animals were doing in a specific experimental setting in a short period of time, there is no reason to consider these values generally descriptive. The same animals might result in using different values in a different setting, and other animals might equally use different values in the same setting. By detaching the parameters from their original definition, we can broaden the scope of the equation and use it to investigate the relationship between the parameters and how they can affect collective behaviour. By doing so I explored how $a$ and $S$ are connected, what their reasonable values for a specific case of environments is, and how different combinations of them lead to different behaviours. Our intention here is not to describe a specific system, but to see how a taxa of vertebrates that follow this collective process respond to their environment by evolving their sociality. In order to do that, I assume separation of time-scales where $a$ remains constant while $S$ changes. In other words, the agents navigate an environment of constant $a$, and over-time adapt to it by employing the values of $S$ that provides the best fitness within it. Mapping the relationship between $a$ and $S$ allows us to explore the most probable routes, as well as the constraints of that evolutionary process. For here on, I will be considering that $a \in (0,1)$ and that $S > 1$: this is the case where the environment is reliable, and the agents are social.

The final aspect of the model relates this Bayesian estimation to the choice the focal agent makes. Following Pérez-Escudero and de Polavieja, 2011, I assume that the focal agent will choose option $A$ with probability $P(A|C, I)$ (that is, with a probability equal to the probability of that option being the best one), rather than making a threshold decision. This assumption allows for the observable reality that an animal confronted with apparently identical conditions and social information may nonetheless make different decisions on different occasions. However, it also introduces a second deviation from rationality, since the probability of choosing correctly is maximised by choosing whichever option has greater than 50% probability of being correct.

Arganda et al., 2012 proposed a unified framework to describe the use of social information by animals, from which stems a variety of different decision rules observed in animals. Four such decision-making rules are considered here, each being a different variation of the same model. Each model variation considers a different way of observing or utilising the same available social information – with the concepts of ‘observing’ and ‘utilising’ being interchangeable here, since I do not assume to know the internal process of the agents.

The dependencies rule considers the full ordered sequence of prior decisions. It constitutes the fuller version of the model in which the focal decision-makers have access to detailed dependent information. The aggregate rule is the first simplification of the original model, as it considers the same social information, only without assuming that the previous decisions are dependent. The aggregate rule is proposed as a reasonable simplification both by Arganda et al., 2012 and by Pérez-Escudero and de Polavieja, 2011 when introducing the unified framework for animals decision-making based on previous observations that some species use Weber's law, i.e. the relative difference between available options during decision-making; this is a reasonable rule on the grounds that many agents tend to follow majoritarian decisions [Gómez-Laplaza and Gerlai, 2011; Perna et al., 2012] without significantly compromising on decision-making accuracy.

The next simplified version is that described by the majority rule; this is reasonable on two grounds. First, observing the direction of the majority is simpler (especially in large groups) and therefore faster and more reliable as a source of social information. Second, if previous agents had chosen independently (as assumed in the derivation of Equation 5.1), then the Condorcet Jury Theorem implies that the probability that the majority is correct will grow quickly with the number of observed agents. Lastly, the dynamic rule is also being considered as the final model variation. Mann et al., 2013 proposes this decision-making rule as an alternative to rules such as the aggregate, as it may convey more accurate information given that the order of prior choices is an important factor in decision-making.

# Chapter 5

# Collective Optimality & Evolutionary Stability of Social Interactions Rules

This chapter analyses a commonly used decision-making model, and investigates several plausible social interaction rules and their evolutionary stability. The results of this Chapter have been published as a research paper in *Physical Biology* Sigalou and Mann, 2023.

The rules considered here stem from the assumption that, despite the actual social information of a focal agent in such a model being an ordered sequence, in practise she is more likely to perceive a simplified version of this ordered sequence. Here I consider one of these plausible versions dealt with by Pérez-Escudero and de Polavieja, 2011 (here referred to as the *aggregate* strategy), one dealt with by Mann, 2018 (here referred to as the *dynamic* strategy) and one additional simplified version of the aggregate (here referred to as the *majority* strategy).

This chapter provides insights on the evolution of sociality for groups using these strategies under two different regimes, explores the invasion dynamics between the strategies and the effect of group size in them, and proposes some necessary constraints between the parameters of the model in order for it to be effectively applied in specific cases.

## 5.1 Collective Optimality & Evolutionary Stability

### 5.1.1 Collective Optimality

I consider how animals employing the simpler versions of the decision strategy described in Chapter 4 will perform in terms of accurately choosing the correct option. I assume without loss of generality that $A$ is the correct choice – that is, I take the reward for choosing $A$ to be 1 (in some arbitrary units of utility or fitness) and the utility of $B$ to be zero. Following the model, each agent chooses either $A$ or $B$ in turn according to the following probabilistic rule:

$$P(\text{Choose A}) = \frac{1}{1 + as^{-\Delta n}} \tag{5.1}$$

where the non-social parameter $a$ defines how reliable the environment is, $s$ how strongly the social information is followed, and $\Delta n = n_A - n_B$ is the social information, specifically the number of agents that have chosen option $A$ minus the number of agents that have chosen option $B$.

When Equation 5.1 is applied sequentially on all group members of a group with $N$ agents, the group will be divided between the two options and will be in one out of $N + 1$ possible configurations, corresponding to the number of agents that have chosen $A$, ranging from 0 to $N$; the probability of each possible configuration depends on the values of $a$ and $s$. By summing over the final configurations and their probabilities for a specific set of $a$ and $s$, I construct a measure for the group's collective behaviour, $\mathbb{E}(n_A)$, that shows the average number of agents that chose option $A$:

$$\mathbb{E}(n_A) = \sum_{i=1}^{N} ip_i \tag{5.2}$$

where $i$ is the number of agents on option $A$ and $p_i$ the probability of i agents being on option $A$. Conceptually I assume that the order of the agents in the sequence is a random permutation for any given decision, such that $\mathbb{E}(n_A)$ represents the expected reward a randomly chosen agent can expect to receive if all agents apply the same decision-making rule.

Throughout this paper I consider $a$ (representing the quality of environmental information) to be a fixed quantity that the agents cannot alter, whereas they may choose a value of $s$ to apply. For a given value of $a$, I define the *collectively optimal* value of $s$ to be that which maximises the value of $\mathbb{E}(n_A)$.

### 5.1.2 Evolutionary Stability

Consider the case of identical agents, all of whom make decisions according to a common rule (Equation 5.1), with the same values for parameters $s$ and $a$. Under this condition, one can identify a *collectively optimal strategy* that maximises the reward for all agents as above, by maximising Equation 5.2 with respect to $s$: this is the strategy that if employed by all agents of the group, it would lead to the optimal $\mathbb{E}(n_A)$ for the group. However, such a strategy is not necessarily evolutionarily stable, since it may be exploited by an individual who applies a different value of $s$. To determine an evolutionary stable strategy (ESS), one must determine a value of $s = s_{\text{ESS}}$ such that if all agents employ this value, no agent can gain by changing their value of to $s = s'$.

Here it is important to be precise in how we calculate the effect of an agent varying $s$. In general the expected reward an agent receives for employing a given $s$ will depend on its position in the sequence, but I assume throughout that agents do not choose these positions, but are instead randomly shuffled in each decision. Therefore, in calculating the expected reward for an agent employing a new value $s = s'$ I average over all the positions in the sequence that this agent might find themselves (with equal probability for each).

Consider a population comprised of identical individuals (all using the same value of $s = s_{\text{group}}$), and one average invading agent using $s' = s_{inv} \neq s_{group}$. In this case, the average group member will have an expected probability of making a successful choice of $P(A)_{group} = \mathbb{E}(n_A)_{group}$, as all agents are identical. The average invader has an expected probability of making a successful choice of $P(A)_{inv}$: this is calculated by using Equation 5.2 for each possible places within the sequence the invader can be in, and taking their average. In this case, the evolutionary stable strategy that the group can employ is the one where no other strategy (i.e. no other value of the parameter $s = s'$) can out-perform. The value of $s$ where this is achieved is calculated analytically, by considering a range of $s$ values, and comparing the rewards for the groups and invader for each one: once these become equal, the respective value of $s$ this is occurring for is $s = s_{\text{ESS}}$. As shown in Figure 5.2(a), there is one such value of $s$ for the case where both group and invader are using the probabilistic decision rule of Equation 5.1, and it is an equilibrium point.

## 5.2 Decision Rules

Above I consider the evolutionary stability of a given parameter value $s$, assuming that all agents employ the same underlying decision rule specified in Equation 5.1. However, given that the derivation of this decision rule includes multiple departures from full rationality, I anticipate that this could be vulnerable to invasion by alternative decision rules. In particular, adhering to the basic mathematical form of Equation 5.1, two alternatives present themselves as natural variations. In the first case the focal agent does not observe the aggregate number of previous decisions in favour of $A$ and $B$, but instead only observes (or responds to) the direction of the majority decision. In this case the appropriate decision rule is:

$$P(\text{Choose A}) = \frac{1}{1 + as^{-\text{sign}(\Delta n)}} \tag{5.3}$$

I term this the *majority* decision rule, since the social information is a lower-dimensional simplification of that in Equation 5.1. This rule is reasonable on two grounds. First, observing the direction of the majority is simpler (especially in large groups) and therefore faster and more reliable as a source of social information. Second, if previous agents had chosen independently (as assumed in the derivation of Equation 5.1), then the Condorcet Jury Theorem implies that the probability that the majority is correct will grow quickly with the number of observed agents.

The second variation I consider is that an agent observes only the most recent decision before its own. Here the decision rule is given as:

$$P(\text{Choose A}) = \frac{1}{1 + as^{-d}} \tag{5.4}$$

where $d = 1$ if the most recent choice was $A$ and $d = -1$ if $B$. I term this the *dynamic* decision rule in line with similar usage by Mann et al., 2014 which investigated an analogous model empirically in humbug damselfish. This rule is motivated by the theoretical findings in the supplementary information of Pérez-Escudero and de Polavieja, 2011 and in Mann, 2018 that more recent decisions should be weighted more strongly by an agent able to fully account for the correlations in previous agents' choices, and by the empirical results of Mann et al., 2014 and Kadak and Miller, 2020 which point to both humbug damselfish and zebrafish responding primarily to the most recent choices of conspecifics.

Similarly to section 5.1, I consider a group of agents, where all the agents are using the same decision-making strategy but one average invader who is using one of the other two strategies. Like before, the group and the invader will respectively generate an expected reward $P(A)_{group} = \mathbb{E}(n_A)_{group}$ and $P(A)_{inv}$, only now these rewards will depend both on the employed value of $s$ and the decision-making strategy, as the strategy determines whether Equation 5.1, 5.3 or 5.4 will be used to calculated the probabilistic decision which will in turn generate the values of those probabilities. As before, as long an average invader is able to reach a value of $P(A)_{inv} > P(A)_{group}$ for any value of $s$ while using a strategy different to the one used by the group (e.g. the dynamic instead of the aggregate), the group's strategy is susceptible to invasion. On the other hand, if an invader is not able to outperform the group for any value of $s$ while using a strategy other than the one used by the group, then the group's strategy is stable against the invaders strategy. Moreover, if a strategy employed by the group is not susceptible to invasion by any of the other two strategies for any value of $s$, then that strategy is evolutionary stable.

## 5.3 Results

### 5.3.1 Collectively Optimal Social Behaviour

The performance of a behaviour is measured by the probability of making the correct decision. This depends on the degree of reliability of the environment's information (value of $a$), and the intensity of following the available social information (value of $s$). Figure 5.1 shows how variation in the social parameter $s$ changes the probability of different group outcomes: panel (a) shows the outcome distribution in the case where $a = 0.9$ and the social parameter is relatively weak ($s = 1.5$). In this case agents are more likely to choose $A$ rather than $B$, but intermediate outcomes (those with a roughly equal proportion of agents choosing $A$ and $B$) are highly plausible. The probability that all agents will choose $B$ is very low. The expected proportion of agents choosing $A$ is 0.54092. In panel (b) I show the outcome distribution for the same value of $a$ (implying the same quality of non-social information) but a greater value of the social parameter ($s = 2.3$). In this case we make an interesting observation: although the probability that all agents will choose $A$ has increased, this has been accompanied by an increase in the probability that all agents will choose $B$, with intermediate outcomes being very unlikely. This has decreased the expected proportion of correct decisions to 0.54029. In panel (c) I show the outcome distribution for the same value of $a$ (implying the same quality of non-social information)

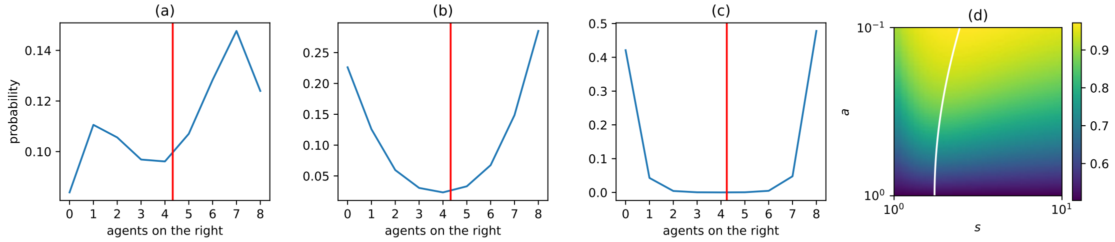

Figure 5.1: Plot (a) shows the probabilities of possible final configurations and the value of $\mathbb{E}(n_A)$ for a group of $n = 8$ agents, $a < 1$ and $s = 1.5$: this corresponds to a case where reliable information is not too strongly followed, leading to very high probability of most agents making a correct decision, and high value of $\mathbb{E}(n_A)$. Plot (b) for a group of $n = 8$ agents, $a < 1$ (same as in plot (a)) and $s = 2.3$: this refers to the case where reliable information is followed strongly, now leading to a decrease in most agents making a correct decision and increase to the probability of most agents making a wrong one, and to a slightly lower value of $\mathbb{E}(n_A)$ compared to when the same information was followed less strongly. Plot (c) for a group of $n = 8$ agents, $a < 1$ (same as in plot (a)) and $s = 10$: this refers to the case where reliable information is followed very strongly, now leading to a further decrease in most agents making a correct decision (with some cases having a probability 0 of occurring) and further increase to the probability of all agents making a wrong one, and to an even lower value of $\mathbb{E}(n_A)$ compared to when the same information was followed less strongly. This demonstrates that over-reliance on social information (even when the environment is reliable) can amplify the potentially wrong decisions made by previous agents. Plot (d) summarises the value of $P_R$ for several combinations of $a$ and $s$, for $a \in$(0,1) and $s \in$(1,10): we observe higher values of $P_A$ that increase as $a$ decreases, while we also observe that across a constant $a$, as $s$ increases $P_A$ decreases, as expected due to the aforementioned cost of over-sociality. For $s > 1$ and $a > 1$ (i.e. social behaviour in unreliable environments) we observe a symmetrically opposite behaviour to being social in reliable environments (around the value $a = 1$.)

but a much greater value of the social parameter ($s = 10$). In this case we notice a further increase in the probability that the agents will choose $A$ or $B$, and a further decrease in the probability of the intermediate outcomes (with most of them having a 0 probability of occurring). This has decreased the expected proportion of correct decisions to 0.5304. In other words, being more social has increased the probability of making a bad decision. This is due to the probabilistic nature of the system: even with reliable non-social information, the agents early in the sequence may still make a poor decision. If the tendency to follow social information is very strong, the improbable but still possible poor decision will be copied by the following agents, resulting in an information cascade, eventually misleading a large proportion of the group. This demonstrates that there is a limit to how strongly social information should be followed to maximise collective accuracy.

The effect of $a$ and $s$ is more widely demonstrated in Figure 5.1(d). This shows the value of $\mathbb{E}(n_A)$ for different combinations these parameters, for $a \in$(0,1) & $s \in$(1,10). We see the observation mentioned previously: being more social in the presence of that information increases the probability of making the best decision up until a point, after which an increase in social behaviour decreases that probability. The white line shown on this panel is the collectively optimal value of $s$ for the corresponding value of $a$.

Based on the calculation of $\mathbb{E}(n_A)$ shown in Figure 5.1(d), it is straightforward to identify the value of $s$ that is collectively optimal, shown by the white line. It is clear that as $a$ increases (i.e. non-social information becomes less reliable), the collectively optimal value of $s$ decreases (agents weight the decisions of others less highly). This makes intuitive sense: as agents are identical, a lower value of $a$ means that other agents are more likely to have made the correct decision, and are therefore more reliable sources of social information.

### 5.3.2 Evolutionary stable strategy

Previous research such as Pérez-Escudero and de Polavieja, 2011 or Farine et al., 2014 has focused on empirical estimation of $a$ and $s$ in Equation 5.1 (or the extended version of this model [Arganda et al., 2012]), but estimating both parameters ignores that in a system under natural selection the values of $a$ and $s$ should be connected so as to optimise the performance of agents' decision-making. Here I will determine this necessary connection between the values of $a$ and $s$ and show that for agents employing Equation 5.1 as a decision-making rule these should not be considered as independent variables.

Above I showed how the collectively optimal value of $s$ varies with the reliability of non-social information, $a$. However, this collectively optimal value of $s$ indicates the value that would be chosen so as to maximise the success of the group as a whole. As noted in the previous section, under individual natural selection such an optimal value cannot be assumed to be stable (resistant to invasion by other strategies). Instead, we must seek an evolutionarily stable value of $s = s_{\text{ESS}}$ such that a group of agents employing this value cannot be outperformed by an individual who changes their value to an alternative $s = s'$; in the following sections, we will be evaluating evolutionary stability via pairwise-invasibility plots [Brännström et al., 2013], i.e. I will be plotting the dynamics between a group whose members all employ the same strategy, and a single invader that is potentially using a strategy different that the group, to assess the invader's invasion success. Figure 5.2 shows the results of this analysis. In panel (a) I show the relative expected rewards for a group employing $s = s_{\text{group}}$ and an invader employing $s = s_{\text{invader}}$ (with non-social parameter $a = 0.9$) – yellow areas show cases where the invader's reward is greater than the rest of the group, and purple vice versa. As the plot shows, there is a single value of $s_{\text{group}}$ (indicated by the red line) such that no invader can profit from choosing a different value. This is therefore the evolutionarily stable value of $s_{\text{ESS}}$ for the particular value of $a$ chosen. Performing this analysis with different non-social parameter values we can map $s_{\text{ESS}}$ as a function of $a$. This is shown in panel (b) (orange line), alongside the previously calculated value of the collectively optimal $s$ (blue line) for comparison. Notably, while both the collectively optimal and ESS values of $s$ show a similar pattern of variation with $a$ (increasing as non-social information becomes more reliable, ie as the value of $a$ becomes smaller), they differ markedly across the range of $a$ values, with the collectively optimal $s$ always being lower than the ESS value. This shows that agents are selfishly motivated to effectively 'use up' the available social information, creating strong correlations with other agents that make their own decisions less useful as a source of information to those that follow them. The collective effect of this is to reduce the average performance of all

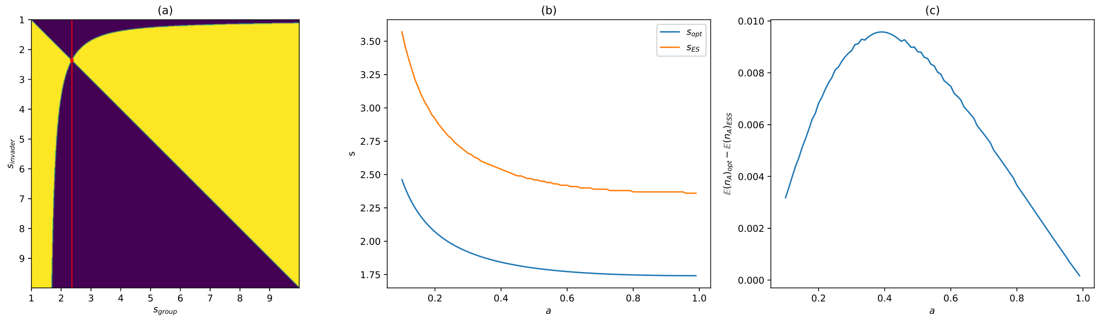

Figure 5.2: Plot (a) shows the dynamics between a homogeneous group and a defector, for different combinations of $s_{\text{group}}$ and $s_{\text{invader}}$ and for one value of $a < 1$: the yellow areas correspond to the case where the invader has a higher probability of making the correct decision, the purple areas to the case where the group has a higher probability of making the correct decision, while the diagonal and the curved line correspond to the cases where they have equal probabilities of making the correct decision (with the diagonal being the special case where they actually have the same behaviour, as it is the line where $s_{\text{group}} = s_{\text{invader}}$). The intersection of the two lines meets at the evolutionary stable point $s_{\text{ESS}}$: notice that while the group remains at that value of $s_{\text{group}} = s_{\text{ESS}}$, for all values of $s_{invader}$ the outcome is that the group will have a higher probability of making the correct decision compared to the invader, thus outperforming her. The group reaches eventually reaches that point due to the existence of invaders: in every other point (for all $s_{\text{group}} \neq s_{\text{ESS}}$ the group is outperformed, and will eventually adopt the invader's $s$ as this is more successful -but once $s = s_{\text{ESS}}$ is reached, no other attempt to invade can be successful. Plot (b) shows the values of $s_{\text{ESS}}$ for the range $a \in [0, 1]$, plotted with the equivalent collectively optimal values $s_{\text{Opt}}$: *collectively optimal* refers to the value of $s$ that the group must use in order to maximise the value of $\mathbb{E}(n_A)$ in the environment it navigates, in the absence of invaders. Notice how the evolutionary stable behaviour is not optimal, but over-social. Plot (c) shows how the difference between $\mathbb{E}(n_A)$ when calculated using $s_{\text{Opt}}$ and using $s_{\text{ESS}}$, i.e. the selective pressure the agents in the group our under due to the invader's presence. We see that the selective pressure is low when uncertainty is low, increases as the uncertainty increases until it reaches a maximum point, and then decreases as uncertainty increases further.

individuals relatively to what they could have achieved had they been able to coordinate on the collectively optimal value of $s$. This 'price of anarchy' [Koutsoupias and Papadimitriou, 2009] (the difference in $\mathbb{E}(n_A)$ under the two strategies) is shown in panel (c) as a function of $a$, showing a peak at $a \simeq 0.4$.

### 5.3.3 Alternative decision rules

Animals may observe and respond to social information in a number of different ways. So far I have considered only one decision-making rule, that assumes agents respond to social information in the form $\Delta n = n_A - n_B$. I now consider the alternative decision-making rules specified in Chapter 4, in terms of the collective behaviour they induce (and whether this is compatible with observations of real animal groups) and their relative performance in decision-making accuracy.

I was able to establish that this framework is compatible with observations, since by using

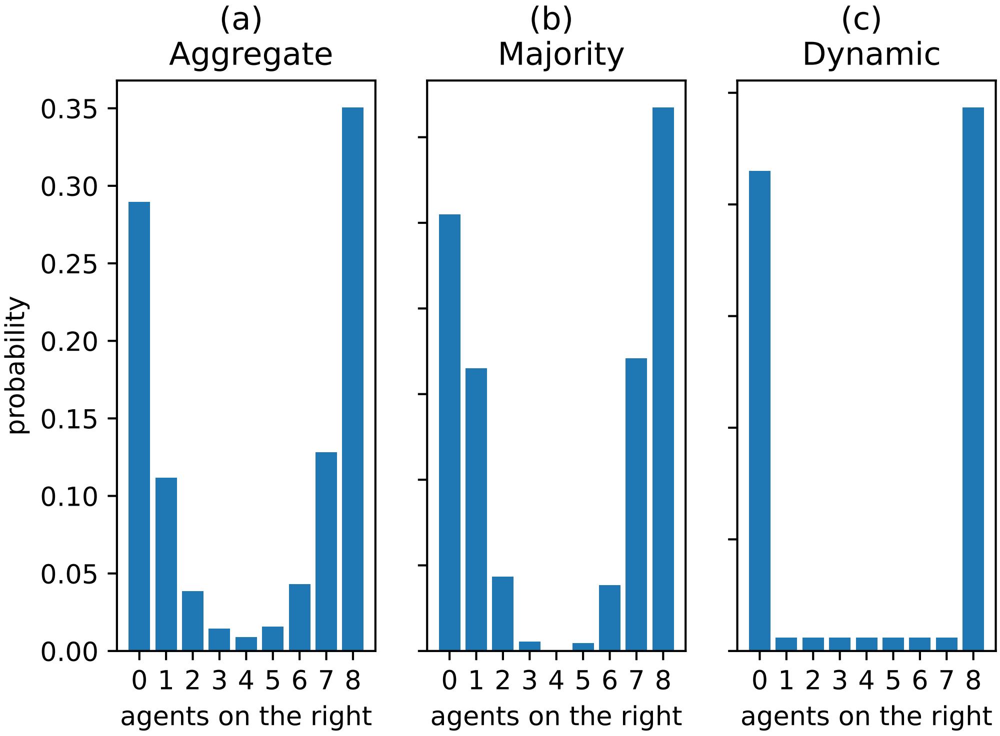

Figure 5.3: Probabilities of final states, for a group of $N = 8$ agents and different strategies. Each plot has a different pair of $a, s$ ($a = 0.9, s = 3$ for panel (a), $a = 0.9, s = 12$ for panel (b) and $a = 0.9, s = 80$ for panel (c)), leading to the same value of $\mathbb{E}(n_A)$ for all strategies and a bias towards consensus-reaching where the probability that all of the agents will make the same choice is high. This demonstrates that we cannot necessarily infer which strategy is being used by the agents simply by noticing that there is a bias towards consensus.

it one can recreate a common characteristic of group decisions, which is the tendency towards consensus decision making, i.e. outcomes in which all agents choose the same option are the most probable [D. J. Sumpter and Pratt, 2009]. All three decision rules I have tested are able to replicate this collective pattern, as shown in Figure 5.3, by selection of appropriate values of $a$ and $s$. This is prima facie evidence that all three models are suitable candidates for modelling collective decisions.

In our analysis above we determined the evolutionarily stable value of $s$ for a group in which agents all employ the decision making rule in Equation 5.1, by analysing whether an invading strategy with a different value of $s$ could outperform the other members of the group. We can extend this stability analysis to ask whether an invading strategy with a different value of $s$ and a different decision making rule can outperform an otherwise homogeneous group.

In Figure 5.4(a) I consider the dynamics between all possible combinations of strategies between group and invader for a group like the one consider before. Along the diagonal are the cases where they both employ the same strategy, while the rest corresponds to the cases where group and invader employ different decision rules. Each column refers to the group using the same strategy (aggregate, majority and dynamic starting from the left), and every row to the defector using the same strategy (aggregate, majority and dynamic starting from the top). The yellow areas signal the cases where the invader outperforms the group, while the purple ones the cases where the group outperforms the invader.

For each decision rule in isolation there exists a single evolutionarily stable value of $s$. Groups employing a single decision rule can therefore be assumed to reach this stable point, where they cannot be outperformed by invaders using different values of $s$. However, they may be outperformed by invaders using a different strategy against it. In each column, the vertical red line signifies the group's evolutionary stable value of $s$. If this line falls within a yellow region for a different decision rule, it signifies that in that case the invader can employ this different decision rule, with the corresponding values of $s$ that are within that region, to outperform the group.

For example, consider the left column of Figure 5.4(a): the top plot shows the equilibrium point for a group employing the aggregate decision rule (Equation 5.1) and the invader's failure to out-compete using the same decision rule (since the vertical line falls exclusively in purple areas). However, if the defector chooses to employ the majority strategy (Equation 5.3), as shown in the middle plot, the same line passes through yellow areas, meaning that there are values of $s$ the defector can employ to outperform the rest of the group. This implies that the aggregate decision rule is not globally stable against invasion by the majority decision rule (assuming that invasions can arise freely on any alternative rule and with any value of $s$, rather than being restricted to local mutations).

This does not imply that the majority decision rule is a stable strategy for the group to employ. Similar inspection of the results as above shows that this rule can be invaded by both the aggregate and dynamic rules. Instead, what this analysis shows is that in a group restricted to employ these three rules, no single decision rule is globally stable. This may eventually lead to the coexistence of different rules in the group, a cyclical transition between rules or the adoption of new rules not tested here. Which of these occurs depends on further assumptions about the evolutionary process beyond the initial invasion, which I do not consider further here. What we can establish from these results is that the aggregate decision rule, which has been widely used as a model for interpreting collective behaviour in real systems, is not stable under the conditions I have described.

So far I have considered a fixed group in which the same agents repeatedly make decisions together; in this case, the assumption is that this group is a population in itself, i.e. evolutionary processes take place in the context of these interactions. However, many animal groups in which decisions are made are transitory, being drawn from a larger population by (for example) fission-fusion dynamics. In such cases the effect of an invader may be different than in the group considered thus far because each agent in the population encounters the invader more rarely (since the invader is rarely part of any randomly selected subgroup), and thus the majority of an agent's rewards are obtained in interactions solely with the dominant phenotype. In other words, in the latter case the population is larger and agents interact not with the whole population at once, but through smaller

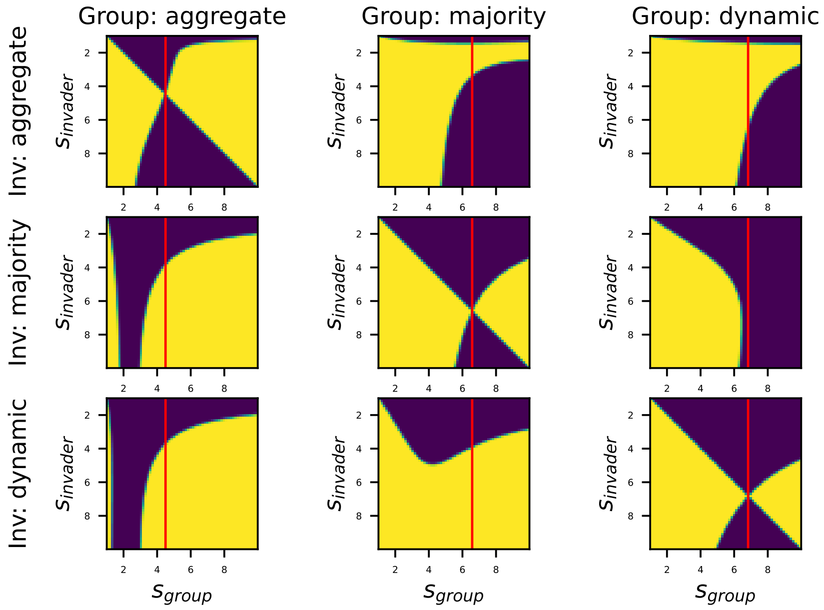

(a) Dynamics in population of fixed size.

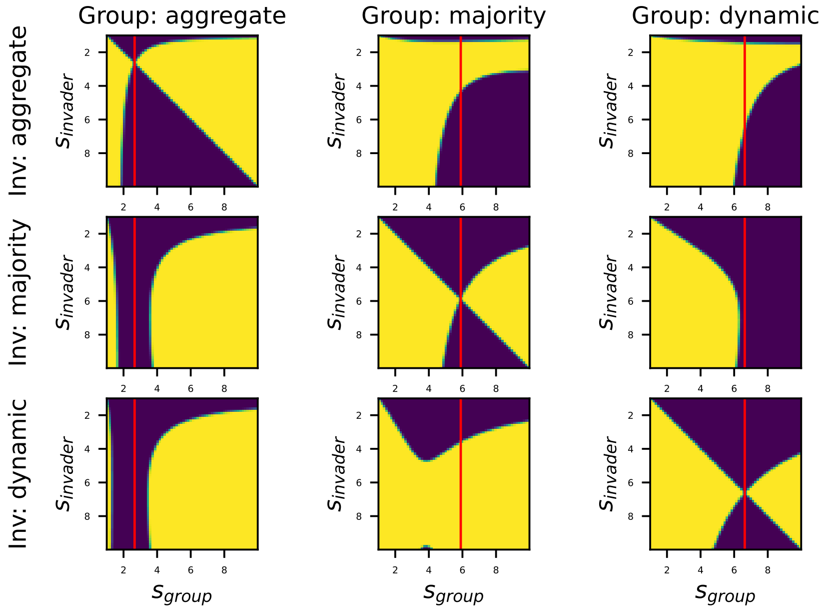

(b) Dynamics in population of infinite size.

Figure 5.4: Group-invader dynamics for different combinations of strategies and $N = 8$, $a = 0.3$.
Plot (a) shows the dynamics for a fixed group: the value of $s_{\text{ESS}}$ for each group's strategy (same across each column) is marked with a red line; if the line passes through yellow areas, the invader's strategy of that row can invade that of the group. For a fixed group, every strategy can be invaded. Plot (b) shows the case of a larger group: here the aggregate strategy cannot be invaded by any other, meaning it is stable.

subgroups. In very large populations the effect of the invader on other agents will be negligible, whereas in a single fixed group the invader may severely disrupt other agents' use of social information across the whole group.

I investigated whether this changed the stability relationships between different decision rules, assuming that in each decision a group of 8 is drawn randomly from an effectively infinite population. In this case we observe different dynamics between population and invader: as shown in Figure 5.4(b), while in this case the simplified and the dynamic strategy can still be invaded from the other two strategies for some values of $s$, the aggregate strategy can't be invaded by either. So in that case, eventually, a larger group will evolve to use the aggregate strategy as this is evolutionary stable, under the assumption that only these three decision rules are available.

## 5.4 Effect of Group Size

So far I have considered a group of size $N = 8$ in an environment with uncertainty of $a = 0.3$. Both group size and uncertainty are important factors in decision-making since group size affects the amount of available social information, and uncertainty impacts the reliability of available social information. In this subsection, five additional regimes are considered to explore the effect that environmental uncertainty $a$ and group size $N$ have in the those dynamics. Figures 5.5, 5.6 and 5.7 show the dynamics between the three strategies for an environment with uncertainty of $a = 0.3$ and increasing group size (of $N = 3, 5, 8$ respectively), while Figures 5.8, 5.9 and 5.4 show the dynamics between the three strategies for $a = 0.9$ for the same group sizes.

Figure 5.5 shows the dynamics between a group of $N = 3$ agents in environment of $a = 0.3$, and a single invader employing a different strategy/value of $s$. In the case of a fixed group all three strategies can be invaded by at least one other strategy but for a very limited range of $s$ values; specifically the aggregate can be invaded only by the majority, the majority only by the dynamic, and the dynamic by both the aggregate and the majority. Hence a fixed group can oscillate between the three strategies. In the case of the infinite group, the aggregate strategy is the only one that cannot be invaded for any value of $s$, while the other two can be invaded in a limited range of $s$ by the aggregate, and in an extremely small range by each-other.

Figure 5.6 shows the dynamics between a group of $N = 5$ agents in environment of $a = 0.3$ and a single invader employing a different strategy/value of $s$. In the case of the fixed group, all strategies can be invaded by at least one more: the aggregate and the the majority can be invaded by both other ones, and the dynamic can be invaded by the aggregate meaning that the group will again oscillate between employing each of the strategies. In the case of an infinite group the aggregate is again the only stable one, while the two other ones can be invaded: the majority by both other ones, and the dynamic by the aggregate.

Here, compared to the case of a smaller group in the same environment $a = 0.3$, we notice that in the case of the fixed group we see an increase in the regions that the aggregate strategy can be invaded. We also notice that the two other strategies can be invaded for more values of $s$, but in their case we also notice a qualitative difference: while before the simplified could be invaded by the dynamic only, now it can also be invaded by the aggregate. The dynamic, which for $N = 3$ could be invaded by both other strategies, now can only be invaded by the aggregate. As for the infinite group, while the aggregate remains the only stable one, we notice that the simplified can now be invaded for larger ranges of $s$ by the other two strategies, while the dynamic can now only be invaded by the aggregate but no longer by the simplified.

Figure 5.7 shows the dynamics between a group of $N = 8$ agents in environment of $a = 0.3$ and a single invader employing a different strategy/value of $s$. In the case of the fixed group the aggregate and majority strategies can be invaded by both others, while the dynamic strategy only by the aggregate; this is similar to the $N = 5$ case, but the invasion can occur for a wider range of $s$ values. In the case of the infinite group, the aggregate strategy is evolutionarily stable, the majority strategy can be invaded by both others, while the dynamic can only be invaded by the aggregate. For the case of $N = 8$, we have a further increase of the range of $s$ values for which a strategy can be invaded by compared to the smaller group sizes, but no qualitative change in the dynamics in terms of which strategy can invade which.

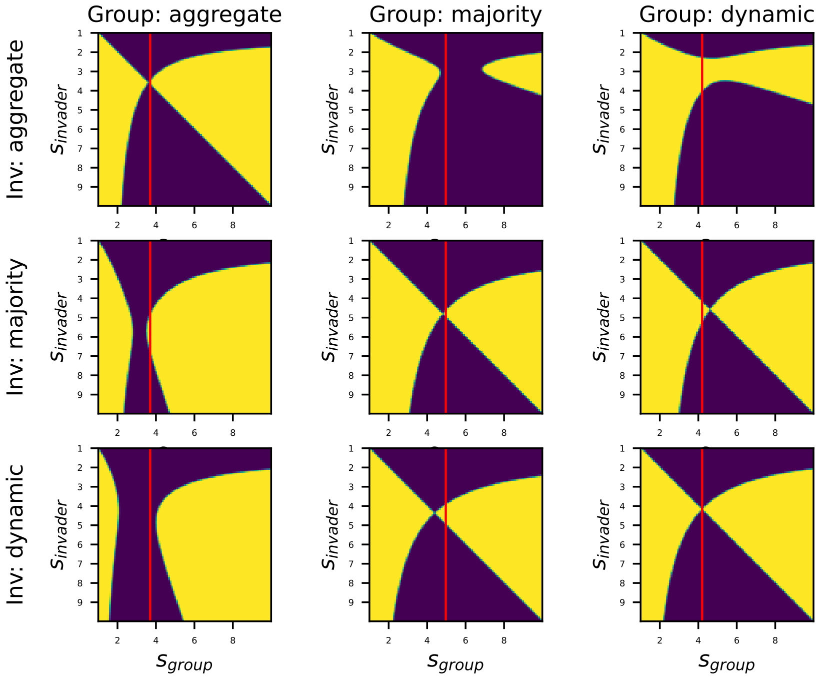

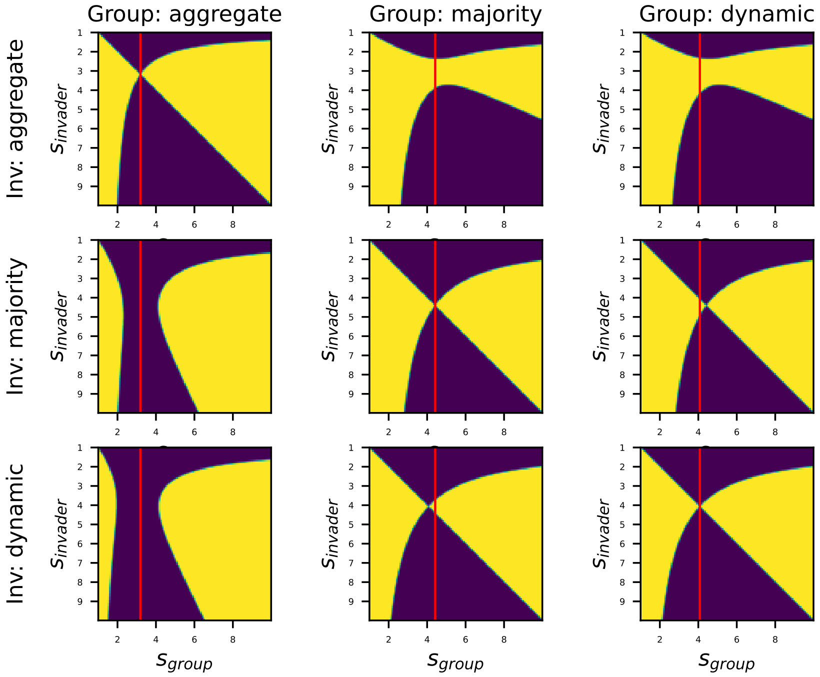

Figure 5.5: Group-invader dynamics for different combinations of strategies and $N = 3$, $a = 0.3$.
Plot (a) shows the dynamics for a fixed group: the value of $s_{\text{ESS}}$ for each group's strategy (same across each column) is marked with a red line; if the line passes through yellow areas, the invader's strategy of that row can invade that of the group. For a fixed group, every strategy can be invaded. Plot (b) shows the case of a larger group: here the aggregate strategy cannot be invaded by any other, meaning it is stable.

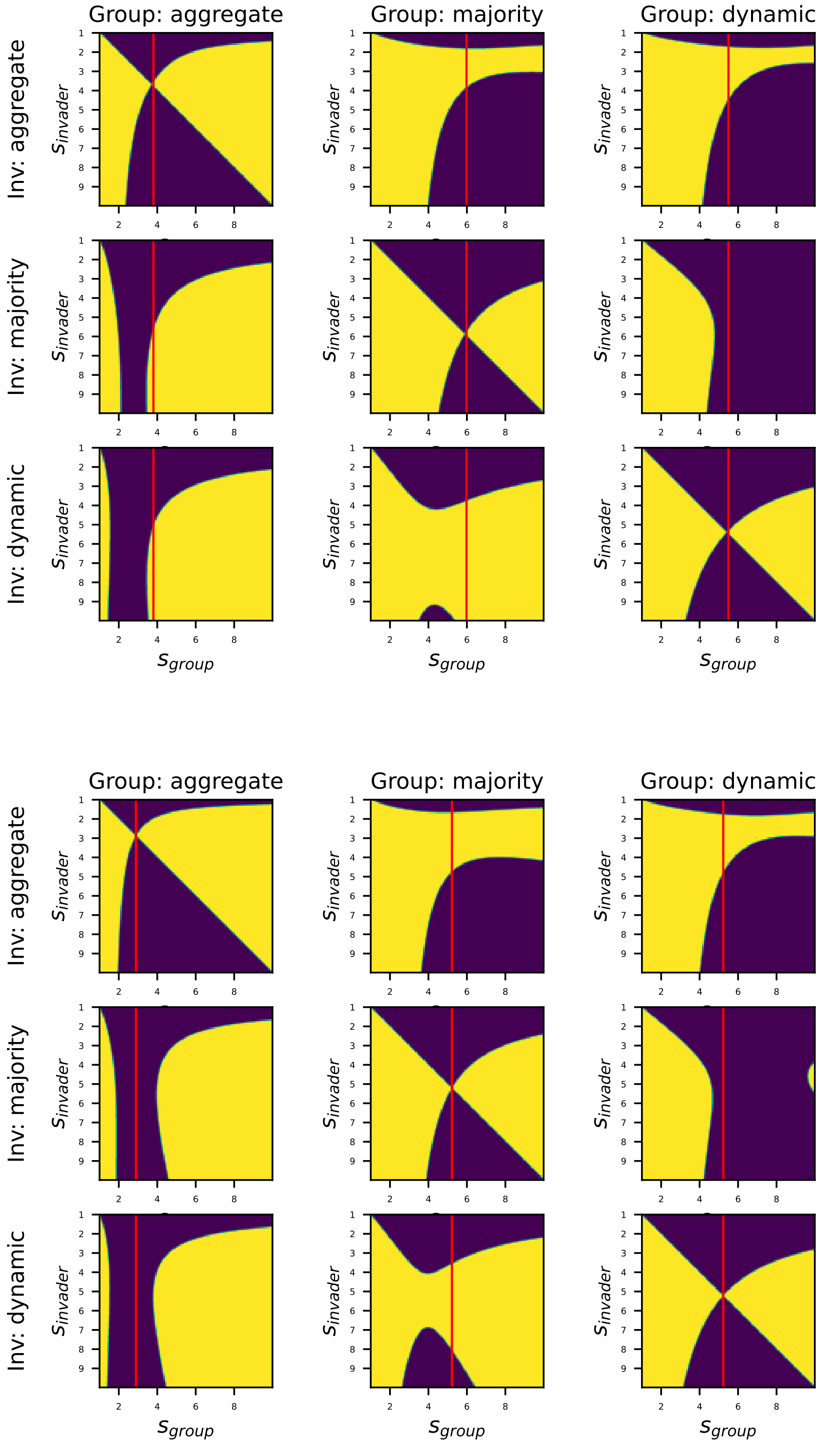

Figure 5.6: dynamics between population and invader, for different combinations of decision-making strategy use where $N = 5$, $a = 0.3$.
Plot (a) shows the dynamics for a fixed group: the value of $s_{\mathrm{ESS}}$ for each group's strategy (same across each column) is marked with a red line; if the line passes through yellow areas, the invader's strategy of that row can invade that of the group. For a fixed group, every strategy can be invaded. Plot (b) shows the case of a larger group: here the aggregate strategy cannot be invaded by any other, meaning it is stable.

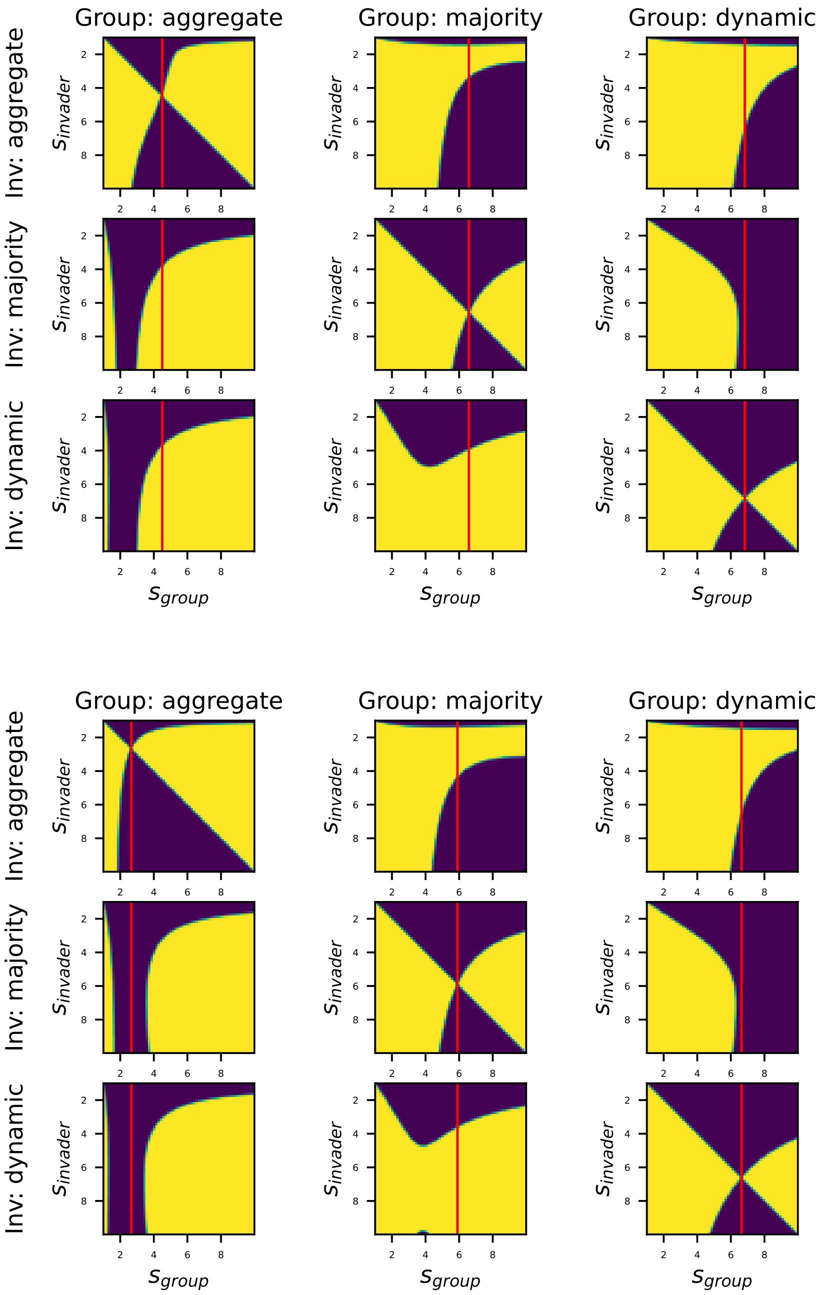

Figure 5.7: dynamics between population and invader, for different combinations of decision-making strategy use where $N = 8$, $a = 0.3$.
Plot (a) shows the dynamics for a fixed group: the value of $s_{\text{ESS}}$ for each group's strategy (same across each column) is marked with a red line; if the line passes through yellow areas, the invader's strategy of that row can invade that of the group. For a fixed group, every strategy can be invaded. Plot (b) shows the case of a larger group: here the aggregate strategy cannot be invaded by any other, meaning it is stable.

Figure 5.8 shows the dynamics between a group of $N = 3$ agents in environment of $a = 0.9$ and a single invader employing a different strategy/value of $s$. In the case of a fixed group all three strategies can be invaded by one other strategy for a limited but larger (compared to the case with $N = 3, a = 0.3$ shown in Figure 5.5) range of $s$ values: the aggregate can be invaded in a small area by the majority, the majority in a very small area by the dynamic, and the dynamic in a very small area by the majority and a larger area by the aggregate. In the case of the infinite group, the aggregate strategy is again the only evolutionarily stable one, while the other two can be invaded in a limited range of $s$ by the aggregate, and in an even more limited range by each-other. In the case of $N = 3$ agents an increase in environmental uncertainty has noticeable impact on the dynamics in the case of a fixed group, but a negligible one in the case of the infinite group. This is expected, since in groups considered as infinite (that is, large enough) the impact of the invader is not as profound as other agents have very low chances of encountering her.

Figure 5.9 shows the dynamics between a group of $N = 5$ agents in environment of $a = 0.9$ and a single invader employing a different strategy/value of $s$. In the case of the fixed group, all strategies can be invaded by at least one more: the aggregate and the majority can be invaded by both other ones, and the dynamic can be invaded by the aggregate similarly to the case of $N = 5, a = 0.3$ but for different ranges of $s$. In the case of an infinite group, the aggregate is again the only evolutionarily stable one, while the two other ones can be invaded (the majority by both other ones, and the dynamic by the aggregate) same as case for $N = 5, a = 0.3$, but now for a larger range of $s$.

Figure 5.10 shows the dynamics between a group of $N = 8$ agents in environment of $a = 0.9$ and a single invader employing a different strategy/value of $s$. In the case of the fixed group, all strategies can be invaded by at least one more: the aggregate and the majority can be invaded by both other ones, and the dynamic can be invaded by the aggregate similarly to the case of $N = 5, a = 0.9$ but for a wider range of $s$ values. In the case of an infinite group, the aggregate is again the only evolutionarily stable one, while the two other ones can be invaded (the majority by both other ones, and the dynamic by the aggregate) same as case for $N = 5, a = 0.9$, but now for a larger range of $s$.

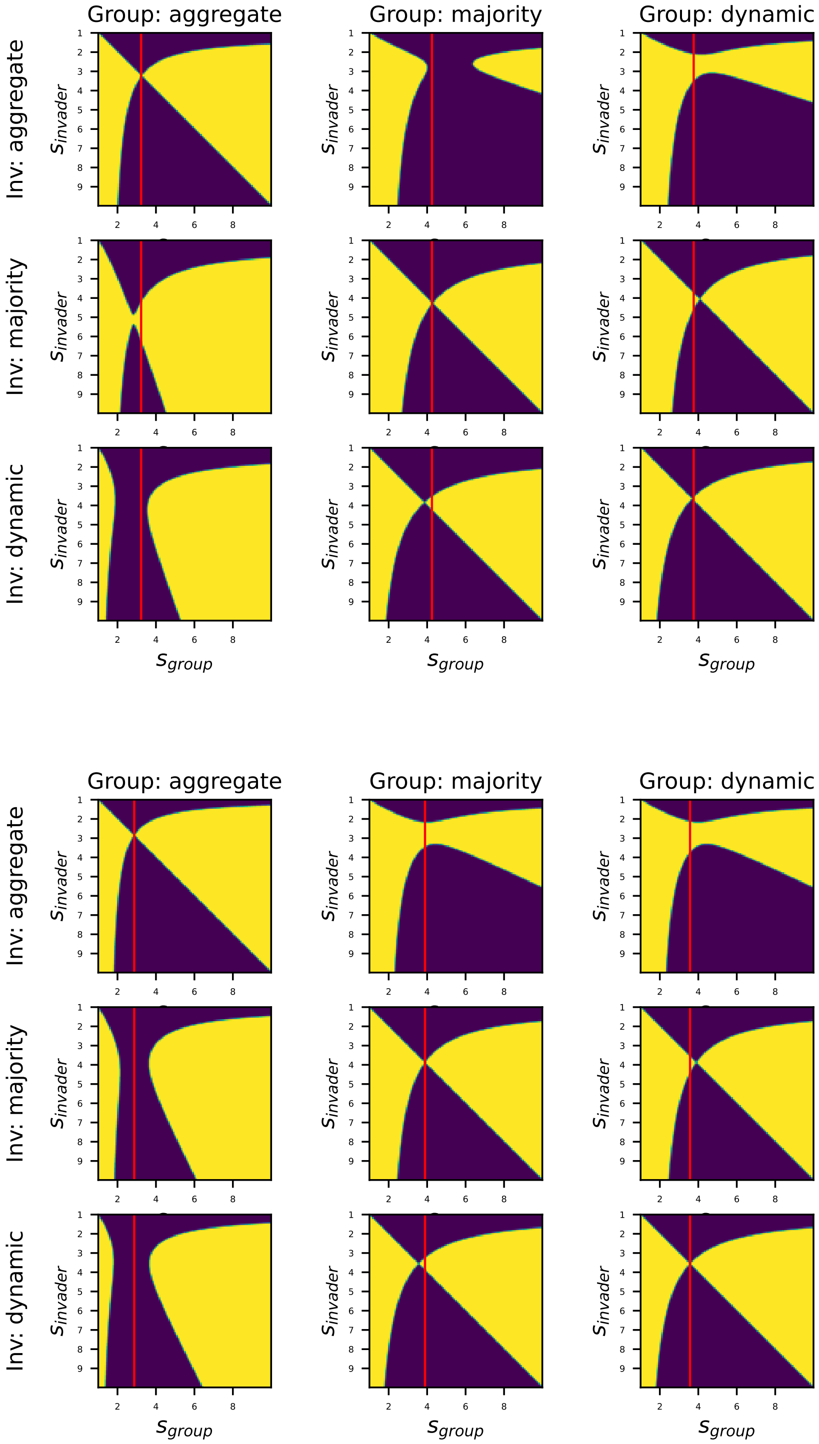

Figure 5.8: dynamics between population and invader, for different combinations of decision-making strategy use where $N = 3$, $a = 0.9$.
Plot (a) shows the dynamics for a fixed group: the value of $s_{\text{ESS}}$ for each group's strategy (same across each column) is marked with a red line; if the line passes through yellow areas, the invader's strategy of that row can invade that of the group. For a fixed group, every strategy can be invaded. Plot (b) shows the case of a larger group: here the aggregate strategy cannot be invaded by any other, meaning it is stable.

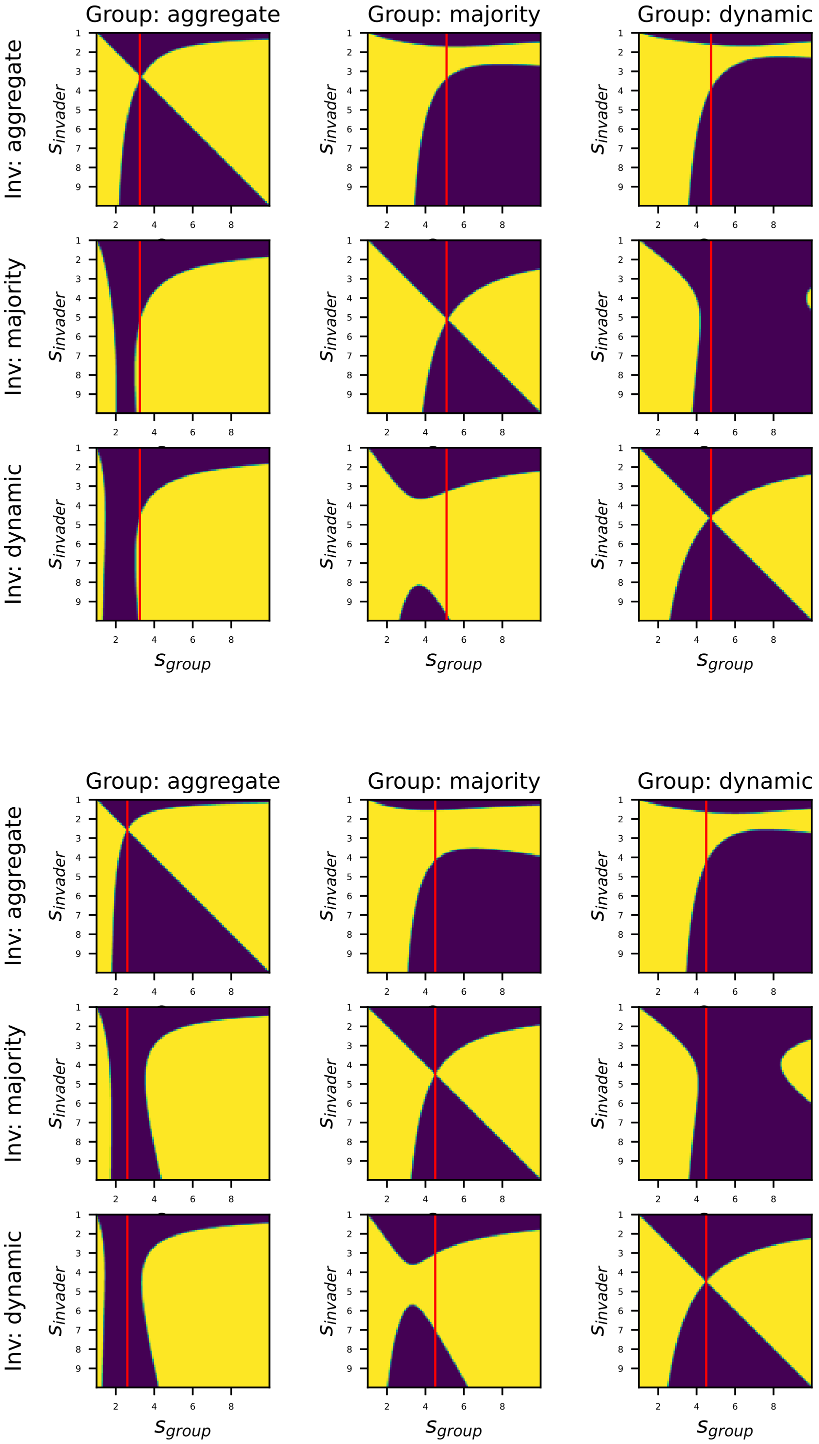

Figure 5.9: dynamics between population and invader, for different combinations of decision-making strategy use where $N = 5$, $a = 0.9$.
Plot (a) shows the dynamics for a fixed group: the value of $s_{\text{ESS}}$ for each group's strategy (same across each column) is marked with a red line; if the line passes through yellow areas, the invader's strategy of that row can invade that of the group. For a fixed group, every strategy can be invaded. Plot (b) shows the case of a larger group: here the aggregate strategy cannot be invaded by any other, meaning it is stable.

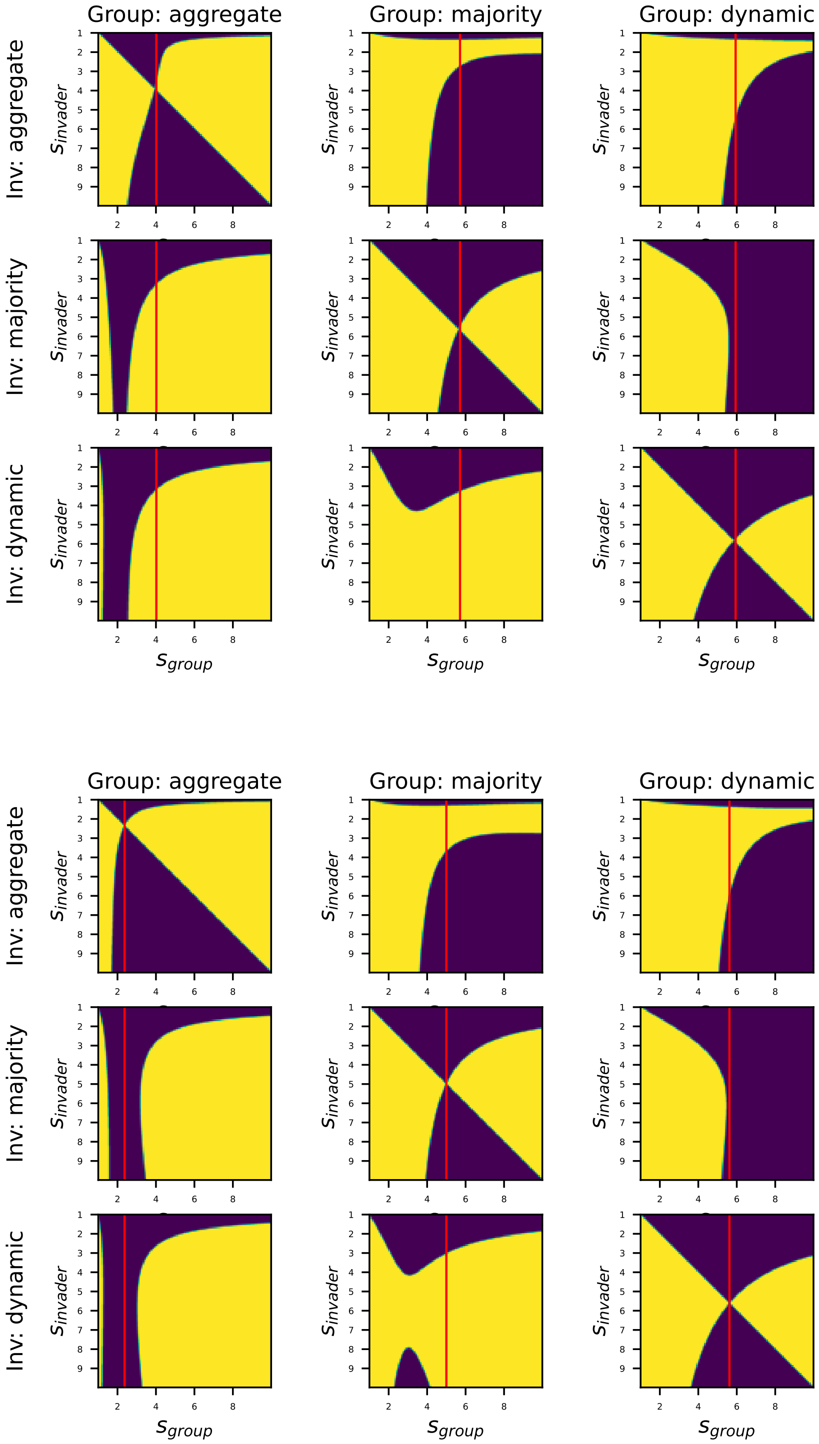

Figure 5.10: dynamics between population and invader, for different combinations of decision-making strategy use where $N = 8$, $a = 0.9$.
Plot (a) shows the dynamics for a fixed group: the value of $s_{\text{ESS}}$ for each group's strategy (same across each column) is marked with a red line; if the line passes through yellow areas, the invader's strategy of that row can invade that of the group. For a fixed group, every strategy can be invaded. Plot (b) shows the case of a larger group: here the aggregate strategy cannot be invaded by any other, meaning it is stable.

For a fixed value of $a$, increasing group sizes lead to an increase in the areas (range of values of $s$, and in some cases strategies themselves) in which the strategies can be invaded in. This is due to the fact that larger group sizes allow for the different effect of each strategy to be observed: a group of $N = 3$ agents does not have enough agents to accumulate bias so a comparison with the other strategies does not yield a very noticeable result. Yet for larger group sizes this effect takes place and now the different strategies lead to different collective outcomes for the groups using them, with the effect increasing with group size (leading to the observed increase in ranges of $s$ values where the strategies can be invaded). The effect of group size is observed in the case of infinite group sizes as well, as in this case an increase in group size equally increases the range of $s$ values that the simplified and dynamic strategies can be invaded; it does not however affect the fact that the aggregate is the only stable strategy among the three. The effect of group size is observed equally for low ($a = 0.3$) and high ($a = 0.9$) value of environmental uncertainty, while it is moreover observed that an increased environmental uncertainty further increases the range of $s$ for which a strategy can be invaded.

## 5.5 Discussion

Collective decision-making emerges from the individuals' decisions, which in turn are affected by the available personal and social information. Here, I use a probabilistic sequential decision-making model, to understand how information affects the quality of decision-making. Since the reliability of social information is tied to the reliability of personal information (in other words, environmental uncertainty), the effect of social information on decision-making depends both on the environmental uncertainty, and level of sociality of the agents.

Then, using the individual decision-making rule I construct a measure that describes the average group behaviour, and study how social behaviour evolves in groups comprised of heterogeneous agents, that share the goal of optimising this collective measure rather than their own probability of choosing the best option.

The probability of an agent making a good decision depends both on the reliability of the available information, and on how strongly this is followed. While in general a more certain environment means that following social information is advantageous, the probabilistic nature and the non-linear form of the decision-making rule mean that an increase in social following is not always beneficial. This is because even in fairly certain environments wrong decisions are possible, and an increase in social behaviour will increase the probability of making a wrong choice, as the agent risks following a misleading cue.

This is also the case for a group made-up from cooperative agents, who do not only try to maximise their own fitness, but also try to provide reliable information to their peers. Given the ambiguity of available information, the sociality of the group members is constrained by the risk of amplifying unreliable information and creating a poor information cascade. For that reason, the collectively optimal behaviour is fairly moderate in intensity, even in environments with very reliable information.

Yet while being over-social is not optimal, it still occurs as a behaviour and I even find it to be evolutionary stable. Previous research has characterised this trait as a failure on the part of the group to behave optimally [Torney et al., 2015], but our evolutionary approach shows that this is in fact the adaptive behaviour of agents attempting to counter the effect of defectors.

I also test whether there is such a thing as an evolved decision-making strategy: I find that this depends on group size. None of the three strategies considered in this paper become evolutionary stable when the group is of fixed size; rather, the group goes through a loop between all the available strategies. For larger groups the group does evolve to use a single strategy, specifically the aggregate. I additionally explore the effect of group size and that of environmental uncertainty on the dynamics between strategies: I find that across all of the treatments considered it holds that for finite size groups no strategy is stable in an environment where the aggregate, the majority and the dynamic are present. In the case of infinite group sizes, regardless of the group size $N$ only the aggregate strategy can become stable in cases where it, the simplified and the dynamic are present. Both the group size and the level of uncertainty are important in determining for what range of values a strategy can be invaded in.

This research is one of the first to consider the evolution of decision-making strategies. As such, I have chosen to make some specific simplifications that may have resulted in simplified results. For example, I find that a small group does not evolve to use one strategy, while a large enough group does; yet I do not consider an important factor: costs. In our model there is no additional cost to using a more computationally or time consuming strategy like the aggregate, instead of a simpler one like the dynamic. I expect that adding costs to our model will potentially change how group-invader dynamics play out.

I have also chosen to define group behaviour as the averaged behaviour of the group-members. While this is informative as it led to insights about behaviour evolution, this approach ignores the nuance of information within a group: depending on the rank, agents have access to different information, and it's sensible to expect this to contribute to different behaviours for agents, depending on their place in the group. Furthermore, I expect this to affect the dynamics between population and invader, since an invader's success depends on what place within the group she manages to place herself. I believe that by including this level of complexity, we will be able to add some insight into the existing literature that researches dynamics between agents with differing (or even conflicting) goals within groups[Conradt and Roper, 2003, 2007; Conradt and Roper, 2009]. This research question is discussed in more depth in Chapter 8.

# Chapter 6

# Limiting Behaviour for Large Groups

The sequential decision-making binary choice process discussed so far, when assumed to consider independence on the order of previous decisions can be approximated by a memoryless stochastic process. This is the case for the last three strategies discussed in Chapter 4, namely the *aggregate*, the *majority* and the *dynamic* strategies. All three strategies only require a knowledge of the system's state at the time of the focal agent's decision, regardless of how the system arrived there. Additionally, after each focal decision, the available social information changes incrementally by $\pm 1$.

In this chapter, the three aforementioned strategies are modelled as discrete Markov processes, and analysed as such to obtain the long-term tendencies of a group employing each one. Following this formulation, it is explored whether – and to what extent – it is possible for a group to shift its preference; for example: can a group employing the aggregate strategy, that has been so far favouring $A$ (i.e. maintained a positive and generally increasing value of $\Delta n$) favour $B$? The following pages provide a preliminary analysis of this question.

## 6.1 Random Walk Analysis of the Aggregate Collective Behaviour Model

Consider the aggregate rule: this is the case where the focal agent observes the difference in numbers of previous choices between the two options. For instance, if 7 agents had chosen option $A$ so far ($n_A = 7$), and 3 agents have chosen option $B$ ($n_B = 3$), the focal agent observes $\Delta n = n_A - n_B = 7 - 3 = 5$. The observed number $\Delta n$ is always an integer. Moreover, as the agents make their decisions it either increases or decreases by 1 (depending on whether each agent chose option $A$ or option $B$).

At each step, the focal agent chooses A with probability $p = \frac{1}{1+as^{-x}}$ and B with probability $q = 1 - p = \frac{as^{-x}}{1+as^{-x}}$, where $x = \Delta n$. If the focal agents chooses A, the system's $x$ value increases by 1, and if she chooses B $x$ decreases by 1.

This process fits the definition of a Random Walk. Let $X_n$ be the value of $\Delta n$ at time $n$; then the process evolves as such:

$$X_{n+1} = \begin{cases} X_n + 1 & \text{with probability } p_{X_n} \\ X_n - 1 & \text{with probability } q_{X_n} \end{cases} \tag{6.1}$$

where $n$ is the step of the process and $X_0 = 0$ as the first agent makes a choice without social information.

As this is a Random Walk, the process is described by the Markov property $X_{n+1} = X_n + Z_{n+1}$ as per the formulation:

$$X_{n+1} = X_0 + \sum_{i=1}^{n} Z_i \tag{6.2}$$

where $X_0 = 0$ is the origin (starting point) of the process, denoting the first agent that makes a decision without social information, and $Z_1, Z_2, ..., Z_n$ are conditionally independent random variables with distribution $\mathbb{P}(Z_i = 1) = p_{X_i}$, $\mathbb{P}(Z_i = -1) = q_{X_i}$. In other words, at the time the first agent is making her decision she will choose $A$ with probability $\frac{1}{1+a}$; depending on whether she choose $A$ or $B$, the second agent will then choose $A$ with $p_{X_i}$ and $B$ with $q_{X_i}$ and so forth.

### 6.1.1 Long-term bias

Over-time, a social group using the aggregate strategy in a biased environment $a < 1$ is expected to amplify this bias. This is demonstrated in the U-shaped curves shown in Chapter 5 such as in Figure 5.3, where depending on the value of $(a, s)$ the agents in the group will be polarised towards one of the choices (instead of having a more even split between them).

In fact, as a group's size increases, this phenomenon snowballs, meaning that once this effect takes over and one choice is favoured, it becomes increasingly difficult for agents to make the opposite choice. This raises the question of whether and when this may happen: here I will show that the properties of the Random Walk indicate that agents in groups using the aggregate strategy tend, at the infinite size limit (i.e. for agents sufficiently back in the sequence), to not choose against the existing social bias. In other words, if an agent observes $n_A = 45$ and $n_B = 2$, she will almost certainly choose $A$.

Following the expression of the process of Equation 6.2, at any given step $t$, the expected value of the process is:

$$\begin{aligned}\mathbb{E}(X_n) &= \mathbb{E}(X_0 + \sum_{i=1}^{n} Z_i)\\ &= \mathbb{E}(X_0) + \sum_{i=1}^{n} \mathbb{E}(Z_i)\\ &= \sum_{i=1}^{n}(p_{X_i} - q_{X_i})\\ &= \sum_{i=1}^{n}(2p_{X_i} - 1)\end{aligned} \tag{6.3}$$

as $\mathbb{E}(X_0) = 0$, and $\mathbb{E}(Z_i) = 1p_{X_i} + (-1)q_{X_i}$

So the value we expect to find the system in at step $n$ depends on the relation between $p_{X_i}$ and $q_{X_i}$, and the step $t$, and specifically on the sign and value of $x = \Delta n$ at this step. If $x > 0$ (social information favours option $A$) and $p_x > q_x$, as $x$ increases then $p_x = \frac{1}{1+as^{-x}}$ increases, approaching 1 as $x \to \infty$. Consequently the sum $\sum_{i=1}^{n}(2p_{X_i} - 1)$ will only take positive values, that also increase with $x$. If $x > 0$ (social information favours option $A$) and $p_x < q_x$, as $x$ increases then $p_x = \frac{1}{1+as^{-x}}$ decreases, approaching 0 as $x \to \infty$. Consequently the sum $\sum_{i=1}^{n}(2p_{X_i} - 1)$ will only take negative values that also decrease with $x$.

On the other hand when $x < 0$ (social information favours option $B$) and $p_x > q_x$, as $x$ decreases then $p_x = \frac{1}{1+as^{-x}}$ decreases, approaching 0 as $x \to -\infty$. Consequently the sum $\sum_{i=1}^{n}(2p_i - 1)$ will only take negative values, that also increase in absolute value with $x$. When $x < 0$ (social information favours option $B$) and $p_x < q_x$, as $x$ decreases then $p_x = \frac{1}{1+as^{-x}}$ increases, approaching 1 as $x \to -\infty$. Consequently the sum $\sum_{i=1}^{n}(2p_{X_i} - 1)$ will only take positive values, that also increase with $x$. This means that for $x \neq 0$, the agents will over time amplify the pre-existing bias towards one of the options, resulting in the expected value $\mathbb{E}(X_n)$ being increasingly further away from the origin as $|x|$ increases and social information aggregates.

This means that for agents further back in the sequence, it becomes increasingly difficult to 'escape' the social bias, as over time it intensifies. Once a group has reached a certain value of $|x|$ – i.e. as the group becomes increasingly polarised towards one of the options – it becomes increasingly difficult for an agent to make the contrary choice (e.g. choose B after a long sequence of prior agents choosing A). It also means that it becomes increasingly difficult for the sign of $x$ to change: so even if a single agent goes against the bias, it's increasingly improbable that this will be repeated long or often enough to reverse the bias towards the non-biased option.

In fact, it is absolutely certain that this will not happen. Imagine a case where a large string of choices towards option $A$ have been made, meaning that at time $n$ we observe $X_n >> 0$: in order for an agent to eventually choose $B$, $X_n$ needs to become $X_n = 0$ at a future time. But given that in a biased Random Walk with $p \neq q$ all states are transient as shown in section 3 this is by definition impossible. Hence, when agents follow the aggregate strategy and it is the case that $p_x \neq q_x$, the social information will overtime remain biased towards one of the options. In the case of the aggregate strategy, it further is the case that all of the agents of a group can make the same choice, as shown in the following subsection.

### 6.1.2 Infinite sequence of same choices

When $p_n \neq q_n$, not only will $E(X_n)$ increase or decrease as $t$ increases, resulting in the vast majority of agents in the group ultimately choosing the same option, but there is a non-zero probability that all the agents will choose the same option.

Let $p_x = \frac{1}{1+as^{-x}}$ be the probability of choosing $A$ – as this is a Random Walk and has the Markov property, then $p_x = P(X_n = x+1|X_{n-1} = x) = 1 - P(X_n = x-1|X_{n-1} = x)$. Then, for an arbitrary starting state $\alpha$ the probability that all agents choose $A$ is $\prod_{x=\alpha}^{\infty} p_x$ for some arbitrary starting point $\alpha$ as the events are disjoint for an arbitrary value of $x = \Delta n$, and:

$$\begin{aligned} \log \prod_{x=\alpha}^{\infty} p_x &= \sum_{x=\alpha}^{\infty} \log p_x \\ &= -\sum_{x=\alpha}^{\infty} \log\left(1 + as^{-x}\right) \end{aligned} \tag{6.4}$$

By the fundamental logarithmic inequality: $1 - \frac{1}{x} \leq \log x \leq x - 1$, so for Equation 6.4 we can say that:

$$\log \prod_{x=\alpha}^{\infty} p_x = -\sum_{x=\alpha}^{\infty} \log(1 + as^{-x}) \geq -\sum_{x=\alpha}^{\infty} as^{-x} = -a \sum_{x=\alpha}^{\infty} \frac{1}{s}^x \tag{6.5}$$

Since only cases where the social information is followed are considered here, it will always be the case that $s > 1 \Rightarrow -1 \leq \frac{1}{s} \leq 1$, and so $\sum_{x=\alpha}^{\infty} (\frac{1}{s})^x = \alpha \frac{1}{s-1}$ is a geometric series. So finally from Equation 6.5:

$$\begin{aligned} \log \prod_{x=\alpha}^{\infty} p_x \geq -a \sum_{x=\alpha}^{\infty} (\frac{1}{s})^x &= -a(\frac{1}{s})^{\alpha} \frac{s}{s-1} \\ \Rightarrow \prod_{x=\alpha}^{\infty} p_x \geq e^{--a(\frac{1}{s})^{\alpha} \frac{s}{s-1}} &> 0 \end{aligned} \tag{6.6}$$

Note that $\prod_{x=\alpha}^{\infty} p_x < 1$ is also the case. So for $a \in$(0,1), $s > 1$ there is a non-zero, but lesser than 1 probability that all agents will make the same choice. This means that an infinite sequence of identical choices can occur in the long-term. This statement is true regardless of our starting point $x = \alpha$, so an infinite sequence of same choices can occur at any time, be it the very first agent of the group or an agent very far along the queue. As per the Markov memoryless property, this can happen regardless of the history of decisions thus far.

## 6.2 Random Walk Analysis of the Majority Collective Behaviour Model

Consider the majority rule: at every step the focal agent observes the sign of $\Delta n = n_A - n_B$, i.e. whether there is a majority on $A$, a minority on $A$, or no difference between $A$ and $B$. The state space here is the same as the one in the case of the aggregate behaviour model, but the agent observes the sign of $x$, i.e. an integer equal to either 1,-1 and 0.

More formally, at each step, the focal agent will choose $A$ with probability $p_x = \frac{1}{1+as^{-sign(x)}}$ and $B$ with probability $q = 1-p = \frac{as^{-sign(x)}}{1+as^{-sign(x)}}$, where $x = \Delta n$. If the focal agents chooses $A$, the system's $X_n$ value increases by 1, and if she chooses $B$ then $x$ decreases by 1. This translates in two distinct cases occurring: the first is the case where the first agent chooses option $A$; once that happens, any subsequent agent will also choose $A$ with probability $p_x$ and $B$ with probability $1 - p_x$. On the other hand, if the first agent chooses $B$ then all following agents will choose $B$ with probability $1 - p_x$ and $A$ with probability $p_x$; in other words, after the initial decision is made the two distinct cases are symmetric, and unless the boundary $X_n = 0$ is crossed it is the case that the process is a simple Random Walk.

By considering then a Random Walk with origin $X_0 = 0$, these probabilities are $p_x = \frac{1}{1+as}$ and $q_x = 1 - p = \frac{as}{1+as}$ when the focal agent is on the origin line, $p_x = \frac{1}{1+as^{-1}}$ and $q_x = 1 - p_x = \frac{as^{-1}}{1+as^{-1}}$ when the focal agent is on the positive space, and $p_x = \frac{1}{1+as}$ and $q_x = 1 - p_x = \frac{as}{1+as}$ when the focal agent is on the negative space, as ultimately the absolute value of $x$ has no effect on the focal's decision – just the sign.

This process fits the definition of a Random Walk, where $X_n$ is the value of $\Delta n$ at time $n$:

$$X_{n+1} = X_0 + \sum_{i=1}^{n} Z_i \tag{6.7}$$

where $X_0 = 0$ as the first agent makes a choice without social information, and $Z_i$ are independent and identically distributed (IID) random variables with $P(Z_i = 1) = p_{X_n}$, $P(Z_i = -1) = q_{X_n}$.

### 6.2.1 Long-term bias

Over-time, a group using the aggregate strategy in a biased environment $a < 1$ is expected to retain this bias, contrary to the aggregate strategy that amplifies it due to the nature of the decision-making rule.

So in this case as a group's size increases, social information does not amplify the pre-existing bias; for both an agent with $i = 2$ and an agent with $i = 50$ the observed social information if there is a majority in option $A$ is going to be the same, namely $x = 1$, despite the actual value of $\Delta n$ being vastly different. Furthermore it is easier (compared to when employing the aggregate strategy) that an agent further down in the sequence observes no social bias towards either of the sides, making it less difficult for agents to make the opposite choice compared to the case where they employ the aggregate strategy.

As this is a Random Walk, the process is described by the Markov property $X_{n+1} = X_n + Z_{n+1}$, and the expected value of the process at step $t$ is $\mathbb{E}(X_n) = n(p_{X_n} - q_{X_n})$. The value we expect to find the system in at step $t$ depends on the step, and on the values of $p, q$ at this step: if $p_{X_n} = q_{X_n}$ then $\mathbb{E}(X_n) = 0$ as there is an equal probability that the focal agent chooses $A$ or $B$. On the other hand if $p_{X_n} \neq q_{X_n}$, as $t$ increases $\mathbb{E}(X_n)$ will over time increase for $p_{X_n} > q_{X_n}$ and decrease for $p < q$. This means that for $p_{X_n} \neq q_{X_n}$, even though the agents don't observe an amplification in bias, the pre-existing bias due to $a$ that is present in the system will lead to an observed preference towards one of the options. Hence for agents further back in the sequence, it becomes increasingly difficult to 'escape' the social bias despite it not being amplified (like in the case of the aggregate strategy). For $p_{X_n} >> q_{X_n}$ or $q_{X_n} >> p_{X_n}$ social following becomes almost deterministic and extreme outcomes can be observed for small groups such as all of the agents choosing the same option (e.g. as demonstrated for a small group in Figure 5.3), but as shown in the next section, in the case of the majority strategy information cascades don't occur in the long-term for large enough groups.

### 6.2.2 Finite sequence of same choices

In the case of the majority rule, information propagation is weaker. Regardless of the level of uncertainty $a$, and the level of sociality $s$, there is a zero probability that all agents will chose the same option.

Let $p_x = \frac{1}{1+as^{-sign(x)}}$ the probability of choosing A. As this is a Random Walk and has the Markov property, then $p_x = P(X_n = x+1 | X_{n-1} = x) = 1 - P(X_n = x-1 | X_{n-1} = x)$. Then the probability that all agents choose A is $\prod_{x=\alpha}^{\infty} p_x$ and:

$$\begin{aligned} \log \prod_{x=\alpha}^{\infty} p_x &= \sum_{x=\alpha}^{\infty} \log p_x \\ &= -\sum_{x=\alpha}^{\infty} \log\left(1 + as^{-sign(x)}\right) \end{aligned} \tag{6.8}$$

By the fundamental logarithmic inequality, $1 - \frac{1}{x} \leq \log x \leq x - 1$, so for Equation 6.4 we can say that:

$$\log \prod_{x=\alpha}^{\infty} p_x = -\sum_{x=\alpha}^{\infty} \log\left(1 + as^{-sign(x)}\right) \geq \sum_{x=\alpha}^{\infty} -as^{-sign(x)} = -a \sum_{x=\alpha}^{\infty} (\frac{1}{s})^{sign(x)} \tag{6.9}$$

In the majority rule, $x = \Delta n$, and the $sign(x)$ can only take three values: 1, -1, 0. When considering the case where everyone is making the same choice, we're considering that either all agents observe $sign(x) = 1$ or $sign(x) = -1$. Then Equation 6.9 will be either $-a \sum_{x=\alpha}^{\infty} (\frac{1}{s})$ or $-a \sum_{x=\alpha}^{\infty} s$.

In either case, the sum $-a\sum_{x=\alpha}^{\infty}(\frac{1}{s})^{sign(x)}=\infty$ and hence doesn't converge, so a statement of the existence of a non-zero probability of everyone making the same choice cannot be made for the majority strategy.

The non-existence of infinite sequences of the same choice can be explained by the weak reinforcement occurring during the majority rule: imagine a large group of agents that have all chosen A. An agent in place $i = 10$ observes the same information as an agent in $i = 2$, as they both observe $sign(x) = 1$. While in the same scenario the agents would become increasingly likely to choose A in the case of the aggregate rule, in this case they are always equally likely to choose A regardless of how strongly it has been favoured thus far. Switching between regions can thus happen in the long-term for the majority rule.

## 6.3 Markov Chain Analysis of the Dynamic Collective Behaviour Model

Consider the dynamic rule: at every step the focal agent observes only the most recent decision. Excluding the very first decision which is made in absence of social information, at every step the system can be in one of two possible states: the previous agent chose A, the previous agent chose B.

From each state, the focal agent can choose either A or B, with complementary probabilities. Then this process is well approximated by a two state Markov Chain (MC) as seen below:

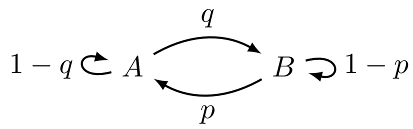

This Markov Chain has transition probabilities $p_{AA} = 1 - q$, $p_{AB} = q$, $p_{BB} = 1 - p$, $p_{BA} = p$ with $p = (1 + as)^{-1}$ and $q = (1 + a^{-1}s)^{-1}$; the transition probabilities can be summarised in the transition matrix $\mathbf{T}$:

$$\mathbf{T} = \begin{bmatrix} 1-q & p \\ q & 1-p \end{bmatrix} \tag{6.10}$$

### 6.3.1 Diagonalisation

To facilitate further analysis, consider the diagonilised expression $\mathbf{T} = \mathbf{Q}^{-1}\mathbf{D}\mathbf{Q}$. The transition matrix $\mathbf{T}$ has eigenvalues $\lambda_1 = 1$ and $\lambda_2 = 1 - p - q$. The eigenvalues can be determined by taking the determinant of $\mathbf{T}$, $det(T - \lambda I)$:

$$\begin{aligned} det(T-\lambda I) = det &\begin{bmatrix} 1-q-\lambda & p \\ q & 1-p-\lambda \end{bmatrix} = 0 \\ \Rightarrow &(1-q-\lambda)(1-p-\lambda) - pq = 0 \end{aligned} \tag{6.11}$$

or

$$(k-\lambda)(1-\lambda)=0 \tag{6.12}$$

where $k = 1-p-q$

Thus the eigenvalues for $\mathbf{T}$ are: $\lambda_1 = 1$ and $\lambda_2 = k$, and we can now proceed with determining the eigenvectors.

Note that in the event that $p, q << 1$, transitions between choosing $A$ and choosing $B$ are very rare, as the focal agent makes the same choice as the previous one almost deterministically. In that case, $\lambda_2 \approx 1$.

The transition matrix $\mathbf{T}$ has eigenvectors $\mathbf{v} = \begin{bmatrix} 1 \\ \frac{q}{p} \end{bmatrix}$, $\mathbf{u} = \begin{bmatrix} 1 \\ -1 \end{bmatrix}$

For $\lambda_1$, the eigenvector $\mathbf{v}$ is found through $(\mathbf{T}-\lambda_1\mathbf{I})\mathbf{v} = 0$:

$$\begin{gathered} \begin{bmatrix} 1-q-1 & p \\ q & 1-p-1 \end{bmatrix} \mathbf{v} = \mathbf{0} \\ \Rightarrow \mathbf{v} = \begin{bmatrix} 1 \\ \frac{q}{p} \end{bmatrix} \end{gathered} \tag{6.13}$$

For $\lambda_2$, the eigenvector $\mathbf{u}$ is $(\mathbf{T}-\lambda_2\mathbf{I})\mathbf{u} = 0$:

$$\begin{gathered} \begin{bmatrix} 1-q-k & p \\ q & 1-p-k \end{bmatrix} \mathbf{u} = \mathbf{0} \\ \Rightarrow \mathbf{u} = \begin{bmatrix} 1 \\ -1 \end{bmatrix} \end{gathered} \tag{6.14}$$

So the transition matrix can be written in the diagonal form $\mathbf{T} = \mathbf{Q}^{-1}\mathbf{D}\mathbf{Q}$, where:

$$\mathbf{D} = \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1-p-q \end{bmatrix} \tag{6.15}$$

$$\mathbf{Q} = \begin{bmatrix} \mathbf{u}, \mathbf{v} \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ \frac{q}{p} & -1 \end{bmatrix} \tag{6.16}$$

And it can be further simplified as $\mathbf{T} = \frac{p}{p+q}\mathbf{QDQ}$, since:

$$\begin{aligned} \mathbf{Q}^{-1} &= \frac{1}{-1-\frac{q}{p}} \begin{bmatrix} -1 & -1 \\ -\frac{q}{p} & 1 \end{bmatrix} \\ &= \frac{1}{1+\frac{q}{p}} \begin{bmatrix} 1 & 1 \\ \frac{q}{p} & -1 \end{bmatrix} \\ &= \frac{p}{p+q}\mathbf{Q} \end{aligned} \tag{6.17}$$

### 6.3.2 Long-term behaviour

When employing the dynamic strategy, the focal agent may only observe the most recent agent but she indirectly has access to older decisions since they are included in the most recent one: the focal agent $i$ observes agent $i-1$, who in turn observed agent $i-2$ and so forth.

Following this, the further 'back' the focal agent is, the more informed she is going to be. Provided that we know the initial probability of the two states A and B, say $\pi_1$, we can find the k-step transition probabilities of agent k.

By the definition of the Markov process, the transition probabilities of the next state only depend on these of the current state:

$$\pi_{i+1} = \mathbf{T}\pi_i \tag{6.18}$$

Here $\pi_i$ define probabilities of choosing either options at time $t = i$: $\pi_i = [p(X_i = A), p(X_i = B)]^T$. Knowing $\pi_1$, the k-step transition probabilities can be written as $\pi_k = \mathbf{T}^{k-1}\pi_1$.

$$\begin{aligned} \pi_2 &= \mathbf{T}\pi_1 \\ \pi_3 = \mathbf{T}\pi_2 = \mathbf{T}\mathbf{T}\pi_1 &= \mathbf{T}^2\pi_1 \\ \pi_4 = \mathbf{T}\pi_3 = \mathbf{T}\mathbf{T}^2\pi_1 &= \mathbf{T}^3\pi_1 \\ &. \\ &. \\ &. \\ \pi_k &= \mathbf{T}^{k-1}\pi_1 \end{aligned} \tag{6.19}$$

The k-step probability can also be written as $\pi_k = \frac{p}{p+q}\mathbf{Q}\mathbf{D}^{k-1}\mathbf{Q}\pi_1$, by replacing $\mathbf{T}$ with its diagonal form.

$$\begin{aligned}\mathbf{T}^{k-1} &= \prod_{k-1} \mathbf{Q}^{-1}\mathbf{D}\mathbf{Q} \\ &= \mathbf{Q}^{-1}\mathbf{D}^{k-1}\mathbf{Q} \\ &= \mathbf{Q}^{-1}\begin{bmatrix}1 & 0 \\ 0 & (1-p-q)^{k-1}\end{bmatrix}\mathbf{Q} \\ &= \frac{p}{p+q}\mathbf{Q}\begin{bmatrix}1 & 0 \\ 0 & (1-p-q)^{k-1}\end{bmatrix}\mathbf{Q}\end{aligned} \tag{6.20}$$

### 6.3.3 Steady-state probabilities

Since this Markov Chain is aperiodic and consists of a single class of recurrent states, the transition probabilities converge to a steady state.

Consider the transition probabilities of the first agent, who makes a decision in absence of any social information: these probabilities then should be $\pi_1 = \begin{bmatrix}\frac{1}{1+a} \\ \frac{a}{1+a}\end{bmatrix} = \frac{1}{1+a}\begin{bmatrix}1 \\ a\end{bmatrix}$. For $\pi_i$ the i-step transition probabilities, consider the sum $\sum_{i=1}^{n}\pi_i$.

The sum $\sum_{i=1}^{n}\pi_i$ denotes the average time spent in the two states up to time $n$. Using the matrix notation, the sum becomes:

$$\begin{aligned}\sum_{i=1}^{n}\pi_i &= \pi_1 + \mathbf{T}\pi_1 + \mathbf{T}^2\pi_1 + \mathbf{T}^3\pi_1 + ... + \mathbf{T}^{n-1}\pi_1 \\ &= (\mathbf{I} + \mathbf{T} + \mathbf{T}^2 + ... + \mathbf{T}^{n-1})\pi_1 \\ &= (\mathbf{I} + \mathbf{Q}^{-1}\mathbf{D}\mathbf{Q} + \mathbf{Q}^{-1}\mathbf{D}^2\mathbf{Q} + ... + \mathbf{Q}^{-1}\mathbf{D}^{n-1}\mathbf{Q})\pi_1 \\ &= \mathbf{Q}^{-1}(\mathbf{I} + \mathbf{D} + \mathbf{D}^2 + ... + \mathbf{D}^{n-1})\mathbf{Q}\pi_1 \\ &= \mathbf{Q}^{-1}\begin{bmatrix}n & 0 \\ 0 & \beta_n\end{bmatrix}\mathbf{Q}\pi_1 \\ &= \frac{p}{p+q}\begin{bmatrix}n+\beta_n\frac{q}{p} & n-\beta_n \\ \frac{q}{p}(n-\beta_n) & n\frac{q}{p}+\beta_n\end{bmatrix}\frac{1}{1+a}\begin{bmatrix}1 \\ a\end{bmatrix}\end{aligned} \tag{6.21}$$

where $\lambda_2 = 1-p-q$, $\beta_n = \frac{1-\lambda_2^n}{1-\lambda_2}$.

Let $\Lambda_A$ and $\Lambda_B$ be the average time spent in A and B up to time $t$ respectively, then these are:

$$\Lambda_A = \frac{p}{p+q}(n + \frac{\beta_n}{1+a}(\frac{q}{p} - a)) \tag{6.22}$$

$$\Lambda_B = \frac{p}{p+q}(n\frac{q}{p} + \frac{\beta_n}{1+a}(a - \frac{q}{p})) \tag{6.23}$$

Then the proportion of time spent in each state will respectively be:

$$\frac{\Lambda_A}{n} = \frac{p}{p+q}(1 + \frac{1}{n}\frac{\beta_n}{1+a}(\frac{q}{p} - a)) \tag{6.24}$$

$$\frac{\Lambda_B}{n} = \frac{p}{p+q}(\frac{q}{p} + \frac{1}{n}\frac{\beta_n}{1+a}(a - \frac{q}{p})) \tag{6.25}$$

The steady state transition probabilities are the values of 6.24 and 6.25 at the infinite limit $t \to \infty$. As $n \to \infty$, $\frac{\beta_n}{n} \to 0$, so these probabilities are:

$$\lim_{n\to\infty} \frac{\Lambda_A}{n} = \frac{p}{p+q} \tag{6.26}$$

$$\lim_{n\to\infty} \frac{\Lambda_B}{n} = \frac{q}{p+q} \tag{6.27}$$

The steady state transition probabilities – i.e. the probability of an agent further down the sequence choosing $A$ or $B$ – depends on the values of $p$ and $q$. This is also the case for the observed pattern of choices: whether there will be long strings of identical decisions, or frequent switches between choosing $A$ and $B$ depends on the values of $p, q$. Since $p$ and $q$ depend on the values of $a$ and $s$, they are constrained to be of similar magnitude; for instance they will either both be very small (that is, $p, q << 1$) or both big rather than $p$ being large and $q$ being small.

This means, that in order for the system to be at an optimal state where option $A$ will mostly be chosen, $p > q$ needs to be the case for a small value of $a$. This is the case because the environmental uncertainty is low making choosing option $A$ more probable, while switching from $A$ to $B$ is more probable than the opposite – and so $A$ will be chosen with a higher probability, switches from $A$ to $B$ will happen with low probability and switches from $B$ to $A$ with a high probability, meaning that we will observe choices on $A$ for the most part, with infrequent brief choices on $B$.

This result produces an interesting implication. If $a$ is small, the difference between $p$ and $q$ intensifies as $s$ becomes larger – in other words, the optimal value of $s$ tends to infinity. This is contradictory to the result of Chapter 5, according to which the value of $s$ is constrained for $a \in$(0,1).

## 6.4 Discussion

The three models considered in this chapter (the aggregate, the majority and the dynamic) are all memoryless versions of the complete model. As such, they are well approximated and analysed as Markov processes. The benefit of this type of analysis lays in the existence of a steady state behaviour: knowing the system and the initial conditions, we know what to expect in the long-run.

I know for example, that in the case of the dynamic model and given an initial non-social decision with $P_A = p = \frac{1}{1+a}$ and $P_B = q = \frac{a}{1+a}$, the probabilities of choosing either side in the long-run are simply $\frac{p}{p+q}$ and $\frac{q}{p+q}$. The amount of time needed for the steady state behaviour to take effect depends by definition on the group size, but the values of these probabilities only depend on the combination of $a$ and $s$ (since $p = (1 + as)^{-1}$ and $q = (1+a^{-1}s)^{-1}$). The first decision is in favour of $A$ with a probability of $(1+a)^{-1}$; that means that for small values of $a$, choosing $A$ is more probable and for larger values of $a$ choosing $B$ is more probable.

Yet since the values of $p$ and $q$ depend on both $a$ and $s$ it is still possible for $p >> q$ even for large values of $a$, in which case eventually most of the agents will choose option $A$, so regardless of what the first agent's decision will be given enough time we expect to see a choice in favour of $A$ with probability $p/(p + q)$. Yet this implies a behaviour for $s$ that appears contradictory to the result of Chapter 5, according to which its value is constrained, since now in order for the optimal outcome to be reached $s$ needs to tend to infinity. This contradiction can be attributed to group size and is further discussed in Chapter 9.

Determining steady state probabilities is not possible for the aggregate and majority strategies given their properties, but it is still possible to predict long-term trends for groups employing them. Additionally, modelling them using the same framework (that of a Random Walk) allows for some direct comparisons between the two.

Focal agents using the aggregate model quickly amplify the existing bias due to $x = \Delta n$, while those using the majority model simply retain it without reinforcing it down the line since $x = \text{sign}\Delta n$. We notice then some important differences: the aggregate model leads to information cascades, as demonstrated by the existence of a non-zero probability of everyone choosing the same option $\prod_{x=\alpha}^{\infty} p_x$ – a trait that does not occur with the majority strategy as this probability cannot be defined since Equation 6.9 doesn't converge. This means that when using the majority model the agents of the group will never over-amplify poor information, but will equally never over-amplify good information. This difference in the impact of bias between the two strategies results in a key difference: the aggregate rule is high-risk, high-reward as it favours extreme outcomes (as all or none of the agents may choose the best option), while the majority model has a fail-safe against poor information cascades but with a catch, as in the long-run it is not possible for all agents to choose the best option.

The analysis here was focused on exploring the progression of decisions for the clear-cut case where the agents were observing information in regions of either $x > 0$ or $x < 0$. This has provided insight into the behaviour of a group under a stable social bias (i.e. social bias indicating only option $A$ or option $B$), while the deviation from this that was explored was that of the first return to the origin $x = 0$ after having spent time in $x \neq 0$. What is not addressed is the behaviour of a group that is transitioning between the regions of $x > 0$ and $x < 0$.

Overall, the preliminary analysis outlined in this chapter provides a basis for a further analysis on the matter; while it strongly indicates that bias amplification leads to information cascades while a lack of it prevents them, what is presented here does not constitute a robust proof. More work is necessary to establish that the presence and amplification of bias is a necessary condition for the existence of information cascades, as well as to understand any potential nuances between the connection between bias and cascades.

# Chapter 7

# Analysis of the Effect of Dependent Social Information

Pérez-Escudero and de Polavieja, 2011 make a simplifying assumption: the observed prior decisions of a focal individual are uncorrelated. This was assumed as their objective was to seek "a model based on probabilistic estimation that can simultaneously give [us] insight into social decision-making and fit experimental data", which they justify since their comparison between the simplified and the complete version shows that the former gives a very good approximation of the latter. It also aligns with the idea that complex collective behaviours can be the result of simple individual rules [Hinz and De Polavieja, 2017; Kadak and Miller, 2020] attributed to one simple common principle [Arganda et al., 2012].

Following this reasoning, Chapter 5 only considers versions of the decision rule that assume uncorrelated decisions. This choice is further justified by considering the fact that many animals, and in many cases, either do not have access to the complete sequence of ordered prior decisions or lack the tools to process them [Mann, 2021].

The idea is that simpler versions of the model are able to well approximate the full version, while being less computationally demanding. This chapter examines this claim to confirm that this is still the case following the necessary constraint imposed in Chapter 5.

## 7.1 Collective Optimality

I consider how animals employing the full version of the decision strategy of Chapter 4 will perform in terms of accurately choosing the correct option. I assume without loss of generality that $A$ is the correct choice – that is, I take the reward for choosing $A$ to be 1 (in some arbitrary units of utility or fitness) and the utility of $B$ to be zero. Following the model, each agent chooses either $A$ or $B$ in turn according to the following probabilistic rule:

$$P(\text{Choose A}) = \frac{1}{1 + aS} \tag{7.1}$$

where the non-social parameter $a$ defines how reliable the environment is, and $S$ is the product of all prior decisions assuming other agents have personal information $\tilde{a}$:

$$S = \prod_{i=0}^{N-1} \frac{P_{b_i,B}}{P_{b_i,A}} \tag{7.2}$$

Note that when parameter $S$ is calculated via Equation 7.2, the focal agent who has personal information $a$ assumed that all previous agents have different social information $\tilde{a} \neq a$ – in other words, she assumes that everyone else in the group has different information from here (better in the case of $\tilde{a} < a$ and worst in the case of $\tilde{a} > a$). This is a similar notion to the parameter $s$ in the case of the three other aforementioned strategies (aggregate, majority and dynamic), since a value of $\tilde{a} < a$ will mean that the focal agent will trust the social information more and a value of $\tilde{a} > a$ will mean that she will trust it less. These two cases are respectively the equivalent of a high and low value of $s$.

When Equation 7.1 is applied sequentially on all group members on a group with $N$ agents, the group will be divided between the two options and will be in one out of $N+1$ possible configurations, corresponding to the number of agents that have chosen $A$, ranging from 0 to $N$; the probability of each possible configuration depends on the values of $a$ and $s$. By summing over the final configurations and their probabilities for a specific set of $a$ and $S$, I construct a measure for the group's collective behaviour, $\mathbb{E}(n_A)$, that shows the average number of agents that chose option $A$:

$$\mathbb{E}(n_A) = \sum_{i=0}^{N} i p_i \tag{7.3}$$

where $i$ is the number of agents on option $A$ and $p_i$ the probability of i agents being on option $A$. Conceptually I assume that the order of the agents in the sequence is a random permutation for any given decision, such that $\mathbb{E}(n_A)$ represents the expected reward a randomly chosen agent can expect to receive if all agents apply the same decision-making rule.

Throughout this chapter I consider $a$ (representing the quality of environmental information) to be a fixed quantity that the agents cannot alter, whereas they may choose a value of $\tilde{a}$ to apply. For a given value of $a$, I define the *collectively optimal* value of $\tilde{a}$ to be that which maximises the value of $\mathbb{E}(n_A)$.

## 7.2 Evolutionary Stability

Consider the case of identical agents, all of whom make decisions according to a common rule (Equation 7.1), with the same values for parameters $\tilde{a}$ and $a$ –that is, every focal agent has personal information $a$ and assumes that everyone else has personal information $\tilde{a}$. Under this condition, one can identify a *collectively optimal strategy* that maximises the reward for all agents as above, by maximising Equation 7.3 with respect to $\tilde{a}$: this is the strategy that if employed by all agents of the group, it would lead to the optimal $\mathbb{E}(n_A)$ for the group. However, such a strategy is not necessarily evolutionarily stable, since it may be exploited by an individual who applies a different value of $\tilde{a}$. To determine an evolutionary stable strategy (ESS), one must determine a value of $\tilde{a} = \tilde{a}_{\text{ESS}}$ such that if all agents employ this value, no agent can gain by changing their value of to $\tilde{a} = \tilde{a}'$.

Here it is important to be precise in how we calculate the effect of an agent varying $\tilde{a}$. In general the expected reward an agent receives for employing a given $\tilde{a}$ will depend on her position in the sequence, but I assume throughout that agents do not choose these positions, but are instead randomly shuffled in each decision. Therefore, in calculating the expected reward for an agent employing a new value $\tilde{a} = \tilde{a}'$ I average over all the positions in the sequence that this agent might find themselves (with equal probability for each).

Consider a population comprised of identical individuals (all using the same value of $\tilde{a} = \tilde{a}_{\text{group}}$), and one average invading agent using $\tilde{a}' = \tilde{a}_{inv} \neq \tilde{a}_{group}$. In this case, the average group member will have an expected probability of making a successful choice of $P(A)_{group} = \mathbb{E}(n_A)_{group}$, as all agents are identical. The average invader has an expected probability of making a successful choice of $P(A)_{inv}$: this is calculated by using Equation 7.3 for each possible places within the sequence the invader can be in, and taking their average. In this case, the evolutionary stable strategy that the group can employ is the one where no other strategy (i.e. no other value of the parameter $\tilde{a} = \tilde{a}'$) can out-perform. The value of $\tilde{a}$ where this is achieved is calculated exact numerically, by considering a range of $\tilde{a}$ values, and comparing the rewards for the groups and invader for each one: once these become equal, the respective value of $\tilde{a}$ this is occurring for is $\tilde{a} = \tilde{a}_{\text{ESS}}$. As shown in Figure 7.2(a), there is one such value of $\tilde{a}$ for the case where both group and invader are using the probabilistic decision rule of Equation 7.1, and it is an equilibrium point.

## 7.3 Results

### 7.3.1 Collectively Optimal Social Behaviour

The performance of a behaviour is measured by the probability of making the correct decision. This depends on the degree of reliability of the environment's information (value of $a$), and the assumptions of others' personal information $\tilde{a}$. Figure 7.1 shows how variation in the social parameter $\tilde{a}$ changes the probability of different group outcomes: panel (a) shows the outcome distribution in the case where $a = 0.5$ and the social parameter assumes better information ($\tilde{a} = 0.25$). The probability that all agents will choose $B$ is very low. The expected proportion of agents choosing $A$ is 0.7442 (or 5.95359 agents).

In panel (b) I show the outcome distribution for the same value of $a$ (implying the same quality of non-social information) and the assumption that others have the same personal information $\tilde{a} = a$. This has decreased the expected proportion of correct decisions to 0.75004 (or 6.0003 agents).

In panel (c) I show the outcome distribution for the same value of $a$ (implying the same quality of non-social information) and the assumption that others have worst personal information ($\tilde{a} = 0.75$). This has decreased the expected proportion of correct decisions to 0.72163 (or 5.77307 agents).

In other words, being more social has increased the probability of making a bad decision. This is due to the probabilistic nature of the system: even with good non-social information, the agents early in the sequence may still make a bad decision. If the tendency to follow social information is very strong, the improbable but still possible bad decision will be copied by the following agents, resulting in an information cascade, eventually misleading a large proportion of the group. This demonstrates that there is a limit to how strongly social information should be followed to maximise collective accuracy.

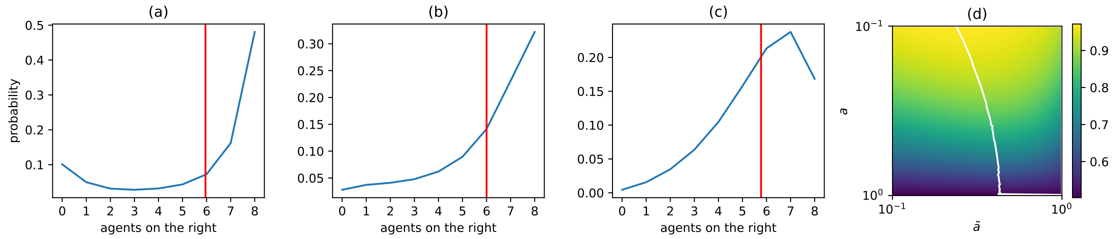

Figure 7.1: Plot (a) shows the probabilities of possible final configurations and the value of $\mathbb{E}(n_A)$ for a group of $n = 8$ agents and $\tilde{a} = 0.5$; this corresponds to a case where others' information is assumed better than the focal's own $a$: Plot (b) for a group of $n = 8$ agents and $\tilde{a} = 0.5$: this corresponds to a case where others' information is assumed equal to than of the focal's own $a$: Plot (c) for a group of $n = 8$ agents and $\tilde{a} = 0.75$: this corresponds to a case where others' information is assumed worst than the focal's own $a$: . Plot (d) summarises the value of $P_A$ for several combinations of $a$ and $\tilde{a}$, for $a, \tilde{a} \in$[0,1]: the white line shows the evolutionary stable value $\tilde{a}_{ESS}$ for the corresponding value of $a$.

The effect of $a$ and $\tilde{a}$ is more widely demonstrated in Figure 7.1(d). This shows the value of $\mathbb{E}(n_A)$ for different combinations of these parameters, for $a, \tilde{a} \in$[0,1]. We notice how for low values of $a$, it's evolutionary stable to assume others have worst information $\tilde{a} > a$ as this minimises the occurrence of poor information loops, while for higher values of $a$ it's stable to assume others have better information $\tilde{a} < a$, as this mitigates against the poor personal information.

Based on the calculation of $\mathbb{E}(n_A)$ shown in Figure 7.1(d), it is straightforward to identify the value of $\tilde{a}$ that is collectively optimal, shown by the white line. It is clear that as $a$ increases (i.e. non-social information becomes less reliable), the collectively optimal value of $\tilde{a}$ also increases (agents assume others have better information than them, and follow them with larger intensity). This makes intuitive sense: as uncertainty increases, agents tend to rely more on social information [Pérez-Escudero and de Polavieja, 2017].

### 7.3.2 Evolutionary Stable Behaviour

Above I showed how the collectively optimal value of $\tilde{a}$ varies with the reliability of non-social information $a$. However, this collectively optimal value of $\tilde{a}$ indicates the value that would be chosen so as to maximise the success of the group as a whole.

Under individual natural selection, such an optimal value cannot be assumed to be stable (i.e. resistant to invasion by other strategies), as an agent employing a different value of $\tilde{a}' \neq \tilde{a}_{group}$ can outperform the group. One must then seek an evolutionarily stable value of $\tilde{a} = \tilde{}_{\mathrm{ESS}}$ such that a group of agents employing this value cannot be outperformed by an individual who changes their value to an alternative $\tilde{a} = \tilde{a}'$; in the following sections, I will be evaluating evolutionary stability via pairwise-invasibility plots [Brännström et al., 2013], i.e. I will be plotting the dynamics between a group whose members all employ the same strategy, and a single invader that is using a strategy different that the group, to assess the invader's invasion success.

Figure 7.2 shows the results of this analysis. In panel (a) I show the relative expected rewards for a group employing $\tilde{a} = \tilde{a}_{\mathrm{group}}$ and an invader employing $\tilde{a} = \tilde{a}_{\mathrm{invader}}$ (with non-social parameter $a = 0.9$): yellow areas denote the cases where the invader's reward is greater than the rest of the group's, and purple vice versa. As the plot shows, there is a single value of $\tilde{a}_{\mathrm{group}}$ (indicated by the red line) such that no invader can profit from choosing a different value. This is therefore the evolutionarily stable value of $\tilde{a}_{\mathrm{ESS}}$.

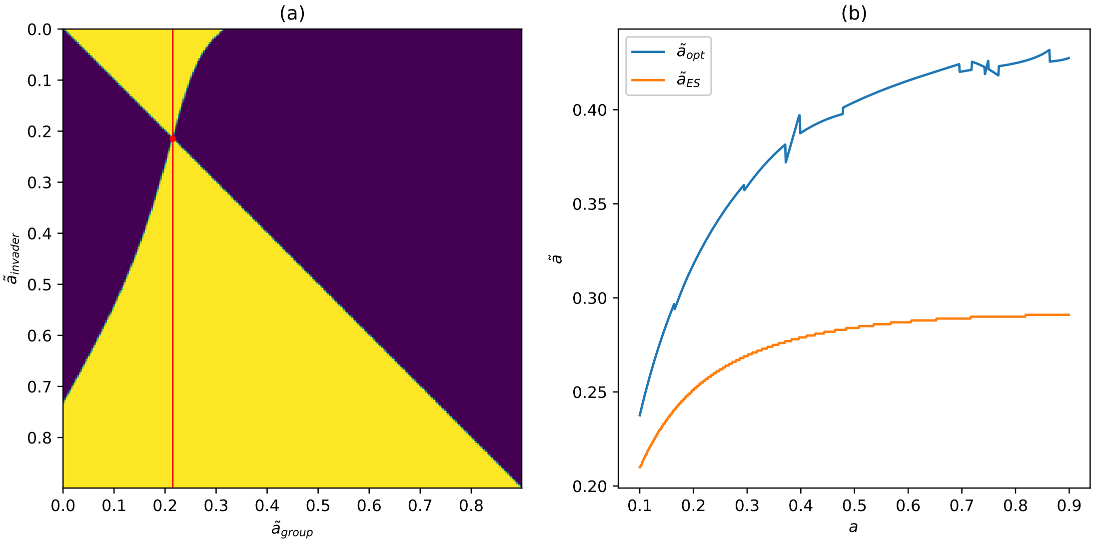

Figure 7.2: Plot (a) shows the dynamics between a homogeneous group and a defector, for different combinations of $\tilde{a}_{\text{group}}$ and $\tilde{a}_{\text{invader}}$ and for one value of $a < 1$: the yellow areas correspond to the case where the invader has a higher probability of making the correct decision, the purple areas to the case where the group has a higher probability of making the correct decision, while the diagonal and the curved line correspond to the cases where they have equal probabilities of making the correct decision (with the diagonal being the special case where they actually have the same behaviour, as it is the line where $\tilde{a}_{\text{group}} = a_{\text{invader}}$). The intersection of the two lines meets at the evolutionary stable point $\tilde{a}_{\text{ESS}}$, denoted by a red vertical line; notice that while the group remains at that value of $\tilde{a}_{\text{group}} = \tilde{a}_{\text{ESS}}$, for all values of $\tilde{a}_{invader}$ the outcome is that the group will have a higher probability of making the correct decision compared to the invader, thus outperforming her. The group eventually reaches that point due to the existence of invaders; in every other point (for all $\tilde{a}_{\text{group}} \neq \tilde{a}_{\text{ESS}}$ the group is outperformed, and will eventually adopt the invader's $\tilde{a}$ as this is more successful – but once $\tilde{a} = \tilde{a}_{\text{ESS}}$ is reached, no other attempt to invade can be successful. Plot (b) shows the values of $\tilde{a}_{\text{ESS}}$ for the range $a \in (0, 1)$, plotted with the equivalent collectively optimal values $\tilde{a}_{\text{Opt}}$: *collectively optimal* refers to the value of $\tilde{a}$ that the group must use in order to maximise the value of $\mathbb{E}(n_A)$ in the environment it navigates, in the absence of invaders. Notice how the evolutionary stable behaviour is not optimal, but over-social.

Performing this analysis with different non-social parameter values allows a mapping of $\tilde{a}_{\text{ESS}}$ as a function of $a$. This is shown in panel (b) (orange line), alongside the previously calculated value of the collectively optimal $\tilde{a}$ (blue line) for comparison. Notably, while both the collectively optimal and ESS values of $\tilde{a}$ show a similar pattern of variation with $a$ (increasing as non-social information becomes less reliable), they differ markedly across the range of $a$ values, with the collectively optimal $\tilde{a}$ always being higher than the ESS value.

This shows that agents are selfishly motivated to effectively 'use up' the available social information, creating strong correlations with other agents that make their own decisions less useful as a source of information to those that follow them. The collective effect of this is to reduce the average performance of all individuals relatively to what they could have achieved had they been able to coordinate on the collectively optimal value of $s$.

## 7.4 Comparison with Other Strategies

While the dependencies strategy is the most complete approach to sequential decision-making, simpler rules are most often used in research as adequate approximations. This is partly done because those rules are simpler to model and analyse; additionally, there is the assumption that a rule such as the dependencies one is too complicated or time-consuming for animals to actually employ.

Here I will compare the dependencies rule against the three other rules considered in this work: the aggregate rule, the majority rule and the dynamic rule. While all rules – including the dependencies one – are considered fairly reasonable in terms of behaviour, an important factor in their plausibility is their efficacy.

All four rules effectively try to reach the best goal: choose the best out of the two available options. They are increasingly complicated as such: dynamic < majority < aggregate < dependencies. The more complex rules are arguably more informed compared to the simpler ones, but they are also more costly (either as they require better sensory organs leading to bigger size and metabolic costs, or because they require more time in order to consider the more detailed information) – so more information-rich is not necessarily better.

Figure 7.3 summarises the comparison between the four rules in the same scenario for a group of 8 agents. Sub-figures (a) and (b) show the comparison between these rules when they are at their collectively optimal state for $a = 0.2$ and $a = 0.7$ respectively; in both cases the dynamic rule under-performs compared to all other rules in terms of the metric $\mathbb{E}(n_A)$, while the $\mathbb{E}(n_A)$ for all other rules coincide. Good environmental information ($a = 0.2$) favours several agents choosing the best option, with the dependencies and the aggregate rules leading to higher probabilities of all agents choosing the best option, and the two rules favouring cases where 5, 6 or 7 (out of 8 total) choose the best option and hence a lower probability that all 8 agents will choose it. Poor environmental information ($a = 0.7$) has a slightly different trend; for the dependencies and the aggregate rule, there is an increase in none or few of the agents choosing the best option, a still higher probability of all or many reaching the best, and intermediate probabilities for the intermediate cases. On the other hand, the simplified and the dynamic rule spike in the intermediate cases, have a lower probability of all (or almost all) of the agents choosing the best option, low probabilities for very few agents choosing it, and zero probability that no agent will choose it. A similar qualitative behaviour is observed for the case where the groups are in their evolutionary stable state, with the only important difference being that the value of $\mathbb{E}(n_A)$ for the simplified rule is between that of the dynamic's and the other two rules.

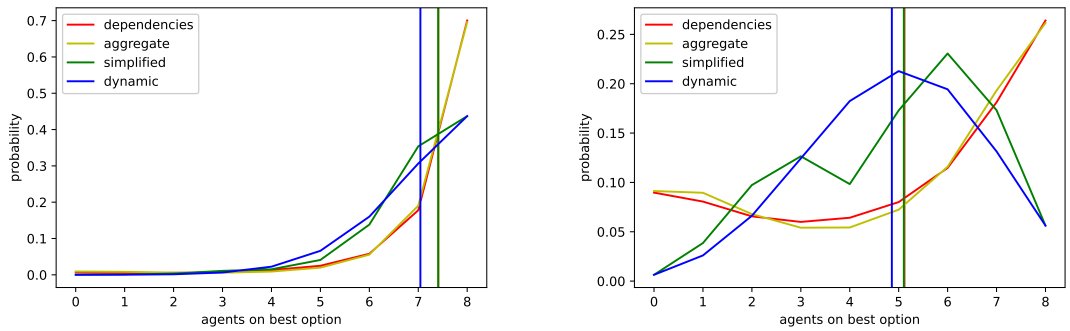

(a) Comparison between the collectively optimal behaviour of the four rules for $a = 0.2$.

(b) Comparison between the collectively optimal behaviour of the four rules for $a = 0.7$.

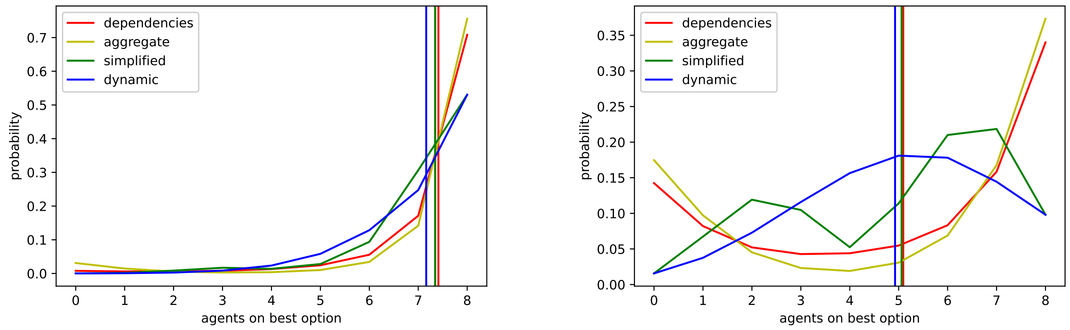

(c) Comparison between the evolutionary stable behaviour of the four rules for $a = 0.2$.

(d) Comparison between the evolutionary stable behaviour of the four rules for $a = 0.7$.

Figure 7.3: The final state probabilities for the optimal social behaviour of each strategy, along with their corresponding $\mathbb{E}(n_A)$.

Above, the four different rules were compared in two ways: on how they affect the collective measure $\mathbb{E}(n_A)$, and on how they affect the final state probabilities distribution.

Effectively, there is no difference between the dependencies, aggregate and majority strategy in regards to $\mathbb{E}(n_A)$: it follows that the necessary chunk of information needed for $\mathbb{E}(n_A)$ to be maximised is whether an option has been majoritarian so far – regardless of to what degree it was favoured over the alternative one, or to whether it was favoured recently or further back in time. There is a more significant difference when comparing for final state distributions, with two 'classes' of rules: 1) those who favour the extreme cases, and 2) those who favour the intermediate cases.

The dependencies and the aggregate rules favour the extremes: under reliable information $a$, the most likely outcome is for all the agents to choose the best option while no-one choosing it is improbable. As the non-social information $a$ becomes less reliable, the extreme case where everyone choose the best option remains the highest overall, but no-one choosing the best option becomes increasingly probable as the information becomes less reliable.

The simplified and the dynamic rules favour the intermediate cases: under reliable information $a$, the most likely outcome is for all the agents to choose the best option while no-one choosing it is improbable; but the extreme case is less probable for these rules compared to the other two ones, while the other cases are more probable compared to the other two rules. As the information $a$ becomes less reliable, the probability that all agents choose the best option ceases to be the most probable and it replaced by more intermediate cases (with the exact case depending on the value of $a$): in that case, the probability that no-one chooses the best option remains equal to zero.

Overall, it doesn't follow from the above comparison that the dependencies rule leads to uniquely better outcomes from all of the other strategies. It does follow though that there are differences in result when considering different levels of detail of social information.

## 7.5 Discussion

The dependencies rule constitutes the complete version of the Bayesian estimation binary choice decision rule presented in Equation 4.1, and it is the most information-dense out of all the models discussed here as it allows the focal agent to observe the exact sequence of previous decisions; so the focal agent knows both what the most popular option is overall, whether there was a sudden shift in preference and how recently, who went against their most recent information e.t.c.

In this version of the model, the baseline assumption about the agents' behaviour is that each agent assumes that all other agents before her have different personal information from her – but the same amongst them. For well-mixed groups employing this rule, that seek to optimise the collective reward $\mathbb{E}(n_A)$ (i.e. the number of agents that choose the best option), there is one collectively optimal value of $\tilde{a}$ that closely depends on the value of $a$ in a way similar to the connection between $s$ and $a$ described in Chapter 5. Furthermore, for well-mixed groups with the above-mentioned collective intention that are infiltrated by an invader who seeks to optimise her own personal outcome (i.e. the probability that she will make the best decision), there is only one evolutionary stable value of $\tilde{a}$ they can employ so that they cannot be outperformed by any $\tilde{a}' \neq \tilde{a}$; this evolutionary stable behaviour is, like the collectively optimal one, closely dependent on $a$.

The evolutionary stable behaviour of a group is always sub-optimal, and specifically over-social. As mentioned previously due to the probabilistic nature of the decision-making model, mistakes are always possible – this is the intuitive explanation behind the constrained values of the collectively optimal behaviour $\tilde{a}_{opt}$. The instability of $\tilde{a}_{opt}$ means that groups develop the tendency to over-rely on social information, making them prone to falling into poor information loops.

Pérez-Escudero and de Polavieja, 2011 argue that simpler versions should be favoured over the complete model, and proposed that in those versions the model parameters $a, s$ are decoupled. In Chapter 5 it is shown that this cannot be the case given the assumptions of the system, as the level of sociality is bound to the environmental uncertainty. Still, there is a basis for considering models that assume weaker dependence on the previous decisions, as this could mitigate against information cascades in the case of poor available information – like the simplified model, which by design does not support infinite identical choices, as demonstrated in Chapter 6.

It follows from the above analysis that simpler models can be reasonably favoured over the complete one, even when taking into account the constraint suggested by Sigalou and Mann, 2023. Despite the dependencies model being the most informative one, it is not necessarily the one leading to the most accurate collective decisions. One reason this might be happening is information overflow: the amount and level of detail it provides cannot be processed by the agents, as their perception tools are limited. Note that although costs are not considered in this project, it makes intuitive sense that this may also be an important factor: complicated strategies are more costly, either because they require more processing time or more advanced cognitive capabilities, so there needs to be a considerable benefit in using them over simpler ones.

From the analysis performed above, there is no evidence to support that. As seen in 7.4, both the average group behaviour as well as they final state distribution are extremely similar for a group employing the dependencies, and the aggregate model, showing no advantage in using the dependencies one over the simpler aggregate model. When looking only at the average group metric, it can be noticed that there is no advantage of using either the dependencies or the aggregate rule instead of the majority one, as they provide an almost identical outcome.

# Chapter 8

# Evolution of Personally Optimal Behaviour

In Sigalou and Mann, 2023 (presented in chapter 5) we consider a well-mixed population: all agents in the group were homogeneous in their sociality, meaning that they evolved to follow social information with the same intensity – i.e. they all shared the same value of $s$ that was either $s_{opt}$ or $s_{ESS}$, depending on the regime. This enabled determining the evolution of sociality during simple processes such as environmental adaptation and the existence of small perturbations in the group, into collectively optimal ($s_{opt}$) and evolutionary stable ($s_{ESS}$) levels respectively.

This assumptions means that the agents have the same utility function, and also have a common goal (optimising $\mathbb{E}(n_A)$). Such conditions are frequently assumed: Pérez-Escudero and de Polavieja, 2011 assumes that the best decision is shared for all agents, and Mann, 2018 assumes that all agents have an identical utility function This leads to a simpler and more solvable model compared to the case where there is such a difference between the agents, however it is limited as differing preferences [Mann, 2020] and conflicts of interest [Conradt and Roper, 2009] are widely spread. The existence of conflicts of interest is of special interest, as they are often framed as obstacles in obtaining cooperative behaviour.

Here the aims are two: 1) to explore the limits of the evolutionary stability as described in Sigalou and Mann, 2023 (also shown in Chapter 5), and 2) to explore under what circumstances cooperation can emerge as a result of an evolutionary process in heterogeneous groups.

## 8.1 Unmixed Groups

When a group isn't well-mixed, the agents assume a fixed position within it rendering the concept of an 'average group member' irrelevant. Instead of an agent being potentially in all possible positions in the queue and hence receiving averaged-out rewards in the long-run, they are now in a fixed position. Given the different available social information for each position in the queue, if all agents in a group share the same level of sociality $s$, each agent now has a different probability of choosing the best option $P_{A_i}$.

For example, for a well-mixed group of 5 agents in an environment of $a = 0.4$ using the aggregate strategy, the collectively optimal level of sociality is $s_{col} = 2.18529$, and on average, each agent $i$ has a probability of $P(A)_i = \mathbb{E}(n_A)$=0.78215 of choosing A over B. For a non-mixed group of 5 agents though, if they continue to employ $s = 2.18529$, then depending on their position, the average probability that each agent $i$ will choose option A is going to be as follows:

| i | $P(A)_i$ |
|---|---|
| 0 | 0.71429 |
| 1 | 0.75622 |
| 2 | 0.79073 |
| 3 | 0.81619 |
| 4 | 0.83333 |

Table 8.1: Average probability of choosing option A, for agent in position i.

Note that each entry of this table is the overall probability of choosing $A$ for each agent. Since this is a probabilistic process, there is more than one possible path available for all agents for $i > 0$; for instance, for a group of $N = 2$ we can either have the following sequences: AA, AB, BA, BB –so agent $i$=1 will choose $A$ either via the path AA or BA. For path AA agent $i$=0 has already chosen option $A$, which happens with probability $p_A = \frac{1}{1+a}$; from that point on agent $i$=1 will choose option $A$ with a probability $p_{AA} = \frac{1}{1+as^{-1}}$. On the other hand, for path BA agent $i$=0 has chosen option $B$, which happens with probability $p_B = \frac{a}{1+a}$, and from that point agent $i$=1 will choose option $A$ with probability $p_{BA} = \frac{1}{1+as^{1}}$. So in the end, the overall probability that agent $i$=1 chooses option $A$ is the weighted sum $p_A p_{AA} + p_B p_{BA}$; for the case of $a = 0.4$ and $s = s_{col} = 2.18529$ these values are $p_A = 0.71429$, $p_B = 0.28571$, $p_{AA} = 0.84528$ and $p_{BA} = 0.53358$ and so $p_A p_{AA} + p_B p_{BA} = 0.152453 + 0.60377 = 0.75622$.

So now, instead of all agents expecting the same long-term outcome of 0.78215, each agent has a position-specific outcome with some under-performing compared to the collective value of $\mathbb{E}(n_A)$, and some performing better. Note that ‘average probability for agent i’ is the arithmetic average of all the possible ways agent $i$ can choose A.

This chapter deals with the evolution of sociality of agents in unmixed groups, where agents optimise their individual outcomes $P(A)_i$ instead of the collective measure $\mathbb{E}(n_A)$. It explores the evolution of personal sociality, and the evolutionary pressure on the agents depending on their position. It also explores the effect that non-cooperators in specifics positions have on the resulting collective measure $\mathbb{E}(n_A)$. The aggregate strategy is being analysed as a baseline, following a supplementary analysis of the majority and dynamic strategies.

## 8.2 Individual Optimality in Collectively Optimal Groups

Being in different positions in the sequence means having access to different social information: the very first agent has no social information, the second agent has social information provided by the first agent, the sixth agent has information from the 5 previous agents e.t.c.

While in well-mixed groups all agents will over time assume every position in the queue (meaning that on average they also personally benefit from the collective $s_{opt}$), this isn't the case for groups that are not well-mixed. In that case, the individual agents have a fixed position i.e. access to a fixed amount of social information.

If we imagine an unmixed group in an environment with fixed $a$, then each agent $i$ within this group can overtime evolve a *personally optimal* behaviour $s_i$, i.e. a level of sociality that maximises their personal probability $P(A)_i$ of choosing option A for their specific position.

### 8.2.1 Being uncooperative when everyone else is cooperative; defector in group

As seen in Chapter 5, the values of $s_{opt}$ are fairly constrained as a result of the probabilistic nature of the system: if sociality is too high, then possible mistakes will be over-amplified so even in environments with low uncertainty (low values of $a$) the evolutionary stable value of $s$ is constrained. As $s_{opt}$ is defined on the basis of optimising average behaviour, one important function is to mitigate against information cascades hence minimising the risk of a large portion of the group choosing the worst option. In a well-mixed group, this is a common concern as agents can assume all positions within the ordered sequence and subsequently need this sort of mechanism.

But when assuming non-mixed groups, this is not the case. The first few agents will never risk receiving over-amplified poor information as they will never be in a position of having access to it; on the other hand, agents in the back of the queue always face this risk. So in an unmixed group, not only do the agents have access to different social information, but also benefit from different levels of sociality. Consider a group of 5 agents, with fixed positions and using the aggregate strategy: agent 0 is the first agent (without access to social information). If agent $i = 1, ..., 5$ is allowed to adapt her $s_i$ value in order to optimise her $P(A)_i$ instead of using the collective $s_{opt}$, she will eventually reach a value of $s_i \neq s_{opt} \neq s_j$, for $j \neq i$, as shown in Figure 8.1. Note that when we consider agent $i$ to be adapting for her personal $s_i$, we assume that all other agents are still using $s_{opt}$.

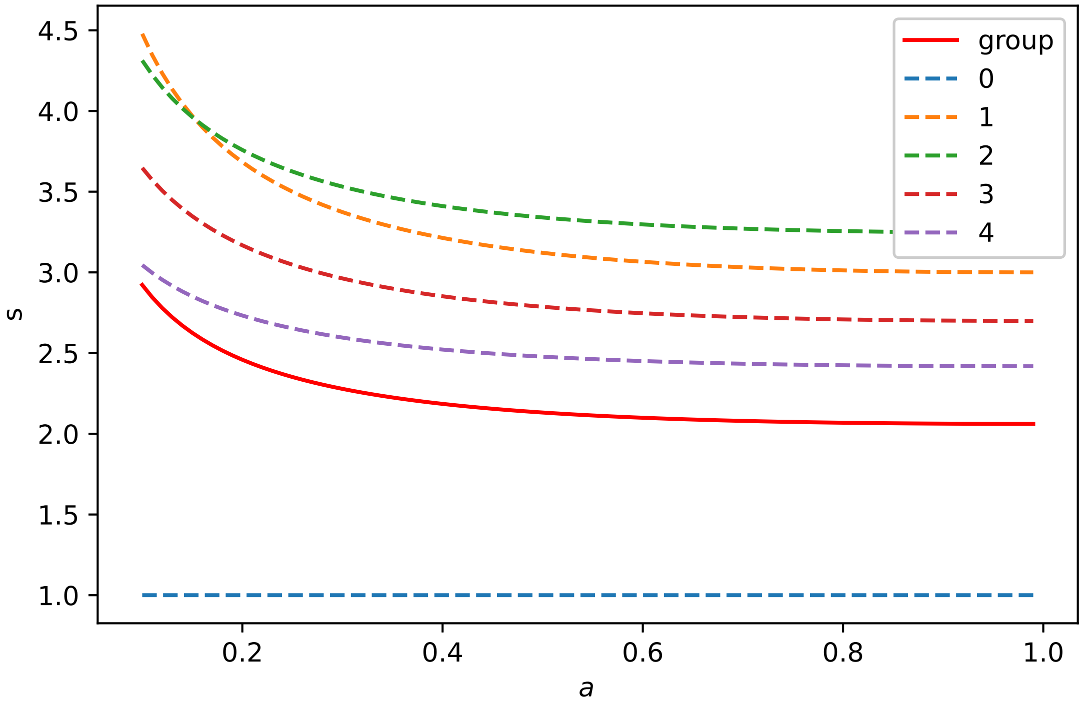

Figure 8.1: Collectively optimal $s$, and individually optimal $s_i$ for all agents in group of $n = 5$ using the aggregate strategy. The first agent (agent 0) has no available social information, hence her $s_i$ is irrelevant and set as equal to 1. All subsequent agents have access to social information provided by agents that are optimising $\mathbb{E}(n_A)$.

Figure 8.1 demonstrates how the collectively optimal behaviour $s_{opt}$ is not personally optimal for any agent with fixed position within the group for the case of the aggregate strategy. For the case of the very first agent (agent $i$=0) who has no access to social information, the value of $s$ is irrelevant and is manually set as equal to 1. In terms of the effect of position on $s_i$ for agents in places $i > 0$, $s_i$ takes lower values as $i$ increases for fixed value of $a < 0.18$. This is due to the fact that the aggregate strategy amplifies bias, and hence more social information also means more bias leading to the benefit of over-sociality to decrease with group position. In more detail: agent $i$=1 has access to only one unbiased decision, and hence benefits from a higher value of $s_i > s_{col}$. Agent $i$=2 has access to two decisions: one independent and one social, hence providing additional but biased information and so her $s_i$ is slightly lower than that of agent $i$=1 for $a < 0.18$. Agent $i$=3 now has information from one independent and two social decisions, so agent $i$=3 is subject to biased social information, so a lower sociality $s_3 < s_2$ will help mitigate against a potential poor information. All agents afterwards are subject to the same issue due to the nature of the aggregate strategy (where bias is amplified fast), meaning that in order to mitigate against accumulated bias as $i$ increases, $s_i$ needs to decrease.

However for values of $a > 0.18$, there is one significant difference: $s_2 > s_1$. This is a highly interesting results, as it is counter-intuitive as the value of $a$ is still very low to provide an obvious explanation to this reversal of the trend. However this could be explained by considering that agent $i$=1 only has access to the social information provided by agent $i$=0 (the agent with no social information, who hence makes a decision purely based on the local value of $a$), making the social information provided to agent $i$=2 not too informative compared to her personal information. Hence, agent $i$=2 evolves to be more social than agent 1 due to this access to more reliable social information.

It follows that for the aggregate strategy, since the agents are now optimising for $P(A)_i$ instead of $\mathbb{E}(n_A)$ they are relying more on social information compared to the case where they were using $s_{opt}$. This leads to the social information provided to the following agents being more amplified, so for higher ranks $i$ the agents are more susceptible to poor information cascades.

| i | $P(A)_i$ |
|---|---|
| 0 | 0.71429 |
| 1 | 0.76023 |
| 2 | 0.79511 |
| 3 | 0.81831 |
| 4 | 0.83392 |

Table 8.2: Average probability of choosing option A, for agent in position i using $s_i$.

It also follows from Figure 8.1 that depending on their position in the queue, different agents are required to compromise to a different degree when optimising for $\mathbb{E}(n_A)$ instead of their own $s_i$. In the case considered above, the first two agents with access to social information (1 and 2) are the ones whose curves are the furthest away from the collective curve, with that difference decreasing the further along an agent $i$ is in the queue. This creates different evolutionary pressure for each agent driving specifics agents to be more uncooperative than others, compared to agents with smaller distances to the collective curve (Figure 8.2).

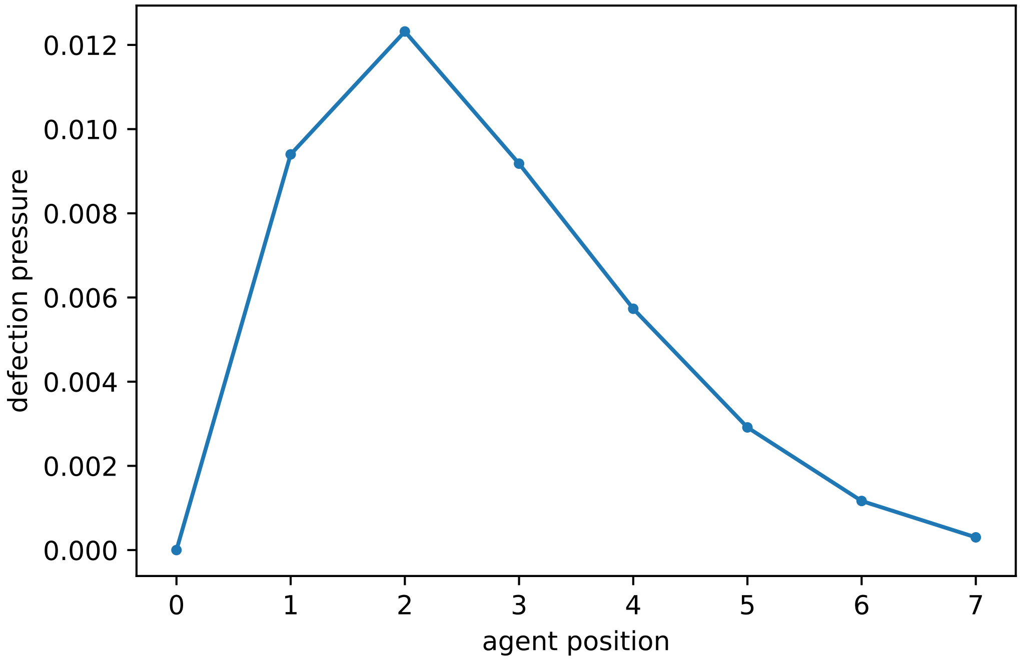

Figure 8.2: The difference between $P(A)_i$ for agent $i$ using $s_i$, minus her $P(A)_i$ using $s_{opt}$, for all $i$ positions in a group of 8 agents with $a = 0.3$ (aggregate strategy). When said difference is positive, it's more profitable for agent i to use $s_i$ instead of $s_{opt}$. Higher positive values indicate a stronger pressure to do so, while lower values a lower one.

### 8.2.2 Defector in specific position

When considering unmixed groups employing the aggregate strategy, a conflict may arise between the agents: the first ones in the queue have an incentive to employ a higher value of $s_i$ which would provide the non-cooperative agents with a better $P(A)_i$ compared to when employing $s_{opt}$. This would consequently provide the following agents with biased information, meaning they would now obtain a lower value of $P_{A_i}$ compared to the case where the previous agents were using lower values of $s_i$. On the other hand, the ones in the back may employ higher values of $s_i$ if the ones further in the front were to have lower sociality $s_i$. If instead the agents all employ the same level of sociality, still the level of 'sacrifice' is different depending on their places in the sequence.

But a conflict may also arise between the potential selfish agent, and the rest of the group. Imagine the same group of 5 agents from before, that employs $s_{opt}$. If an agent instead uses $s_i$, this will not only affect her own outcome, but the final $\mathbb{E}(n_A)$ for this group.

The impact of this effect depends on the uncooperator's position $i$, with the exception of agent $i$=0 who has no social information. Figure 8.3 shows the $P(A)_i$ and $\mathbb{E}(n_A)$ of a group of 5 agents, for the case where agent $i$ is cooperative (uses $s_{opt}$) and the case where she is selfish (uses $s_i$); plot (a) is for a case of low uncertainty $a = 0.4$, while plot (b) is for a case of high uncertainty $a = 0.7$. In both cases, the same trend observed in tables 8.3

and 8.2 is repeated: the agents have a higher $P(A)_i$ when using $s_i$ instead of $s_{opt}$. In terms of how an uncooperative agent impacts the groups, this also depends on her position $i$: when she is one the first few agents with social information, the impact on the group is negative, but otherwise it's non-existent. This happens because the further forward she is in the sequence, the more agents she can impact, while the further back, the fewer agents she can impact.

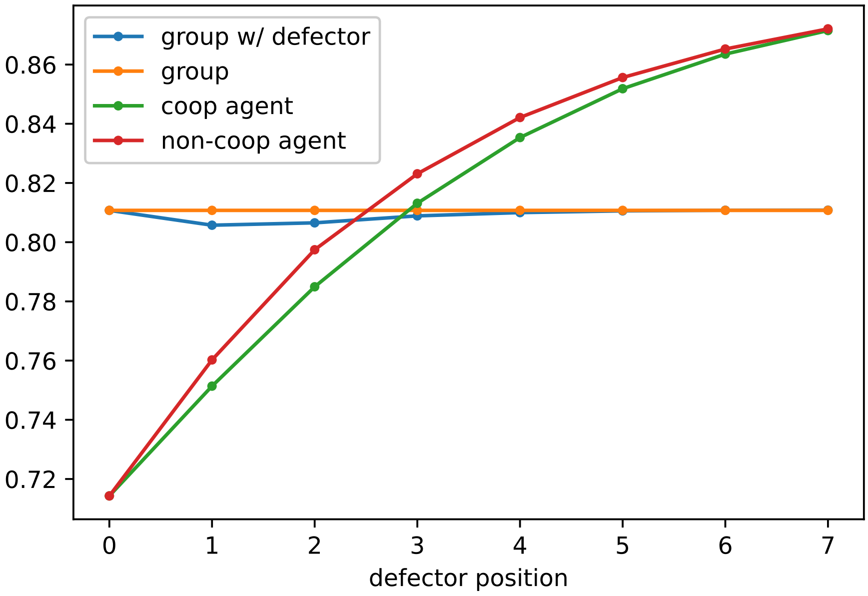

(a) a=0.4

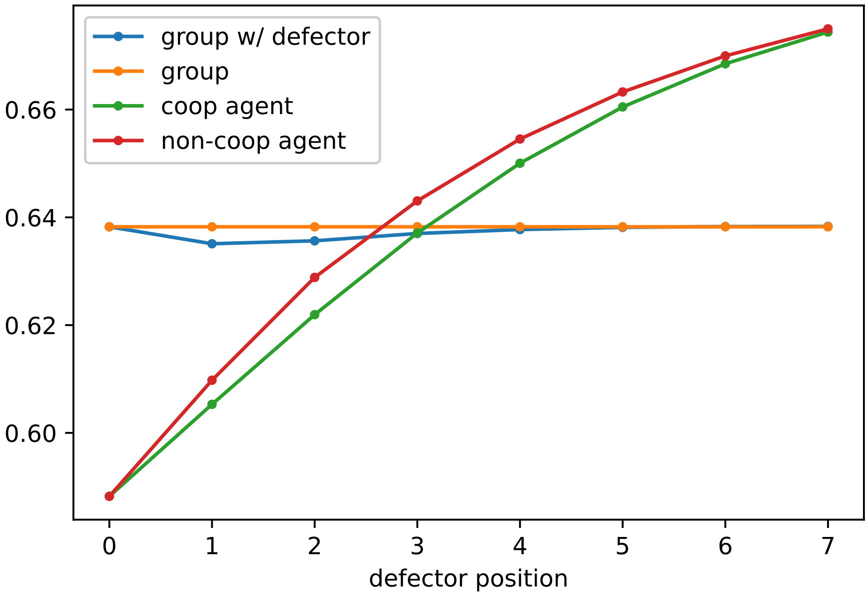

(b) a=0.7

Figure 8.3: Comparing the $P(A)_i$ of a cooperative and uncooperative agent in place $i$. The uncooperative agent over-performs the cooperative agent, with this effect being more noticeable in early positions within the sequence. A group with an uncooperative agent in early positions is negatively impacted: when the uncooperative agents is further back, the group is impacted increasingly less.

This is more clearly demonstrated in Figure 8.4, with the added effect of uncertainty $a$. Plot (a) is the difference between $P(A)_i$ of the uncooperative agent, minus that of a cooperative agent in the same position $i$; positive difference means that not being cooperative is more beneficial for that position under that value of uncertainty $a$; note how in the case of the aggregate strategy, for all positions $i$ and values of $a$ it is either more beneficial to be uncooperative, or leads to the exact same $P(A)_i$ (for cases of low uncertainty $a$ and agents further back in the sequence). Plot (b) is the difference between a group's $\mathbb{E}(n_A)$ with an uncooperative agent in position $i$, minus that of a group with a cooperative agent in that same position; negative values mean that the group on average is negatively impacted by having a selfish agent in position $i$ for that $a$. Note that for the case of the aggregate strategy, the presence of an uncooperative agent does not lead to a positive outcome for the group, as it either leads to lower values of $\mathbb{E}(n_A)$ or to a value of $\mathbb{E}(n_A)$ that is equal to that of a group with no uncooperative agents.

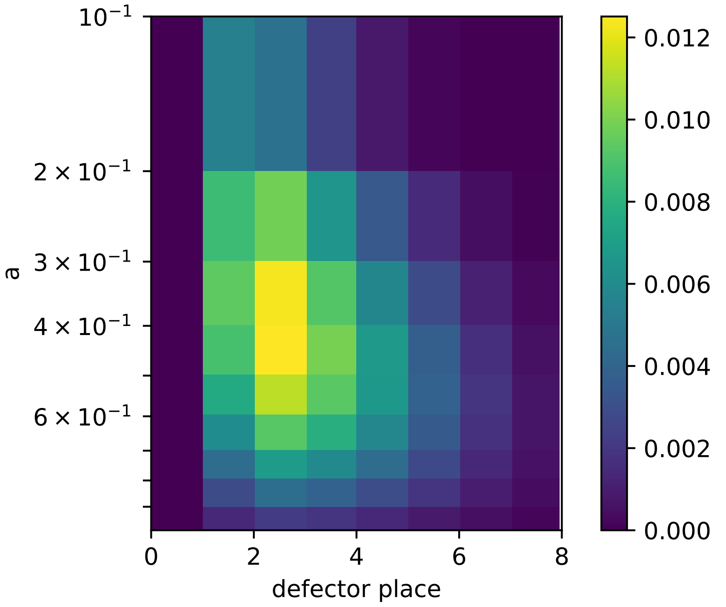

(a) comparing agents

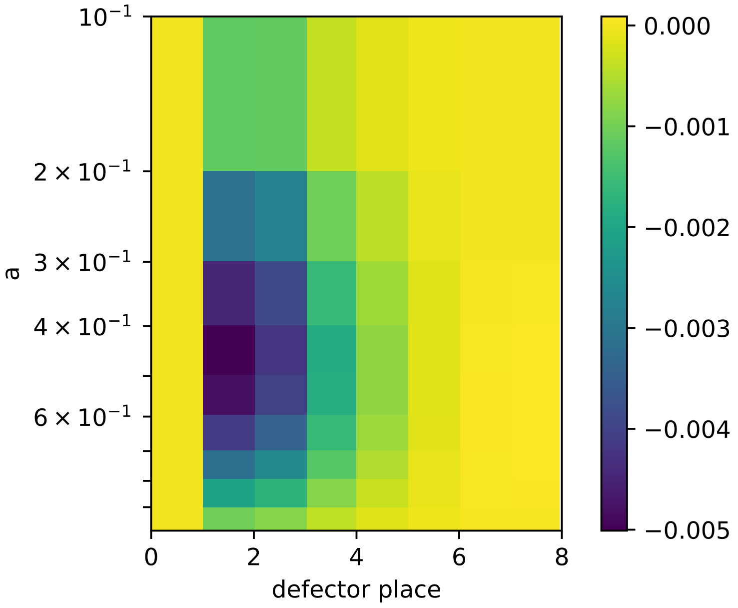

(b) comparing groups

Figure 8.4: Comparison of individual and group success when an uncooperative agent is present in position $i$ (aggregate strategy). (a): difference between uncooperative agent in place d, and cooperative agent in place d for d∈[0,8], in a range of $a \in(0,1)$. The plotted value is the uncooperative agent's $p(A)_i$ minus the cooperative agent's $p(A)_i$ in the same position; zero means it makes no difference to be uncooperative, positive means it pays more to be so.
(b): difference between group with an uncooperative agent in place d for d∈[0,8] and group with only cooperative agents, for a range of $a \in(0,1)$. The plotted value is the $\mathbb{E}(n_A)$ of the group with an uncooperative agent in position $i$, minus the $\mathbb{E}(n_A)$ of the group without an uncooperative agent in that position; zero means it makes no difference whether there is an uncooperative agent, and negative means it's worse for the group to have an uncooperative agent in that position.

Essentially, Figure 8.4 demonstrates how much more beneficial it is for an agent in position $i$ to employ $s_i$ instead of $s_{opt}$ (8.4 (a)), and how much worst it is for the group's $\mathbb{E}(n_A)$ when this happens (8.4 (b)). An agent being uncooperative when at the final few positions has no significant effect in neither the agent or the group. On the other hand, an agent being uncooperative in early positions has an impact on both.

As shown in Figure 8.4(a), the very first agent with access to social information has the highest selective pressure to be uncooperative when the uncertainty (value of $a$) is very low; for increasing values of $a$ the second agent with access to social information is the one with the highest selective pressure to be uncooperative. As discussed earlier in this chapter, this can be explained if we consider that generally the very first agent with access to social information can only observe the decision of the agent without social information; except for the case of very low uncertainty ($a < 0.2$) this decision is not very informative as the first agent with social information can observe the environment in the same way as agent $i$=0. Hence for $a > 0.2$, due to this lack of informative social information for agent $i$=1, it is rather agent $i$=2 that has access to reliable social information worth following with increased intensity. For subsequent agents, the social information incorporates enough bias to lead to lower benefits in being uncooperative.

The impact that an uncooperative agent has on the group's $\mathbb{E}(n_A)$ is almost complimentary to what we observe in Figure 8.4(a): the positions with the highest selective pressure for an agent to be uncooperative in, are leading the group's value of $\mathbb{E}(n_A)$ to decrease. However, the position and uncertainty regime that lead to the highest difference in $P(A)_i$ between uncooperative and cooperative agent (that is, the agent in position $i$=1 for $a < 0.2$) does not coincide with the position and regime (that is, agent in position $i$=2 for $0.2 < a < 0.8$) that has the most effect on the group's $\mathbb{E}(n_A)$, as demonstrated in Figure 8.4(b).

### 8.2.3 Other strategies

Due to the different nature of the three strategies considered across this work, the above results regarding the individual optimality of the agents are not generalisable across them. While the aggregate strategy leads to a swift accumulation of bias due to its formulation (where the social information is $\Delta n$), this is not the case for the simplified and the dynamic strategies. The different formulation of the simplified and the dynamic strategies not only leads to different trends in the values of $s_i$ as $i$ increases, but also in differences in the evolutionary pressure of the agents, as well as in the effect an uncooperative agent has on the group's resulting $\mathbb{E}(n_A)$.

#### Majority Strategy

Contrary to the aggregate, the majority strategy does not accumulate bias, but retains the bias as the agents now only observe the sign of the difference ($\text{sign}(\Delta n)$) between sides. This means that agents far along the sequence do not encounter over-amplified information, and can be more trusting of it compared to the case of the aggregate one. As a consequence, it is beneficial for agents to increasingly trust the available social information, as demonstrated in Figure 8.5, where agents evolve to larger values of $s_i$ compared to the case of the aggregate strategy.

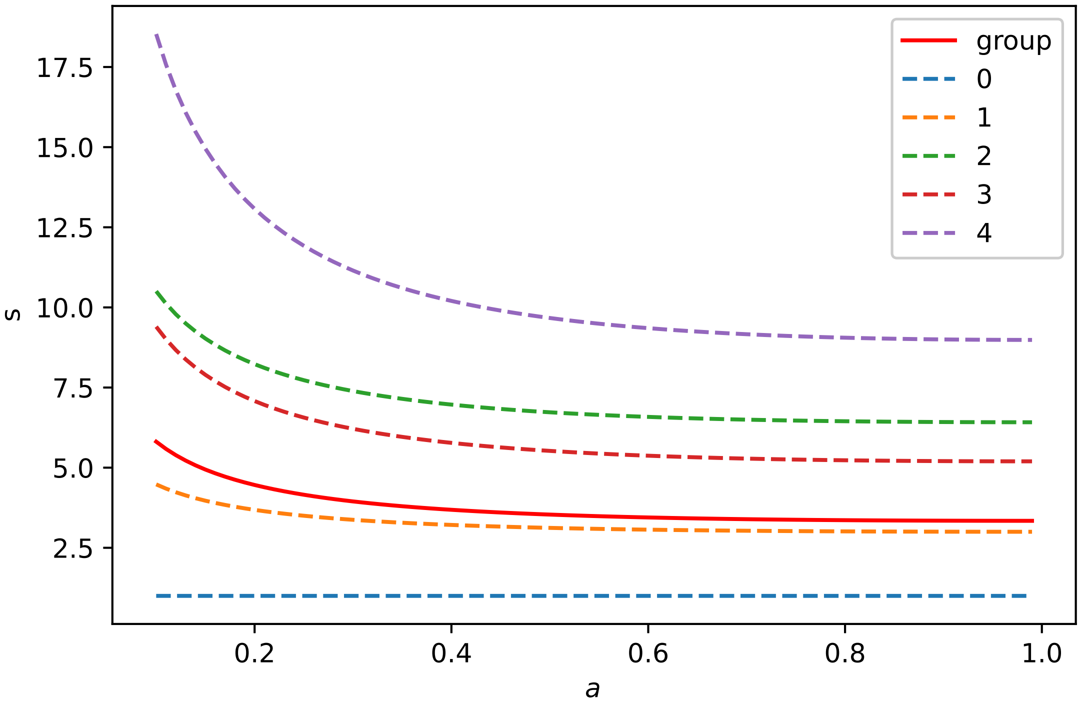

Figure 8.5: Collectively optimal $s$, and individually optimal $s_i$ for all agents in group of $n = 5$ using the simplified strategy. The first agent (agent 0) has no available social information, hence her $s$ is irrelevant. For all the other agents, they all have values of $s_i$ larger than $s_{opt}$.

Figure 8.5 demonstrates how the collectively optimal behaviour $s_{opt}$ is not personally optimal for any agent with fixed position within the group for the case of the simplified strategy. In terms of the effect of position on $s_i$ for agents in places $i \geq 1$, the value of $s_i$ takes alternating higher and lower values; so agent 2 has a higher value than agent 1, agent 3 has a slightly lower value than agent 2 but higher than agent 1, agent 4 has a higher value than agents 2 and 3.

For the case of the very first agent (agent 0) who has no access to social information, the value of $s$ is irrelevant and is manually set to 1. For the case of the first agent with social information (agent $i$=1) who now has access to some limited social information, $s_i$ takes a value lower than the collectively optimal value of $s_{opt}$. Agent $i$=2 has access to social information from two separate agents, one of whom is independent hence it makes sense to follow them strongly with $s_2 > s_1$. Agent $i$=3 has access to some additional (reliable) social information from agent $i$=2. However, an agent in this position is more prone to observing biased information compared to the previous agent (agent $i$=2) who may potentially observe a tie ($sign(\Delta n) = 0$); for this reason, $s_3$ is lower compared to $s_2$. Agent $i$=4 has now access to even more (reliable) social information, and similarly to agent $i$=2 is likely to observe a tie (in other words, to observe at least seemingly unbiased information), hence $s_4 > s_2$.

In the case of the majority strategy, the further back in the sequence an agent is, the higher the evolutionary pressure to be uncooperative (Figure 8.6). This trend is not linear, as agents that cannot observe a tie $sign(\Delta n) = 0$ have a lightly lower evolutionary pressure compared to their exactly previous agent who may observe a tie – a trend similar to the one observed for their personally optimal values $s_i$ demonstrated in Figure 8.5. So overall, the further back an agent is in the sequence, the more beneficial it will be for her $P_{A_i}$ to be uncooperative.

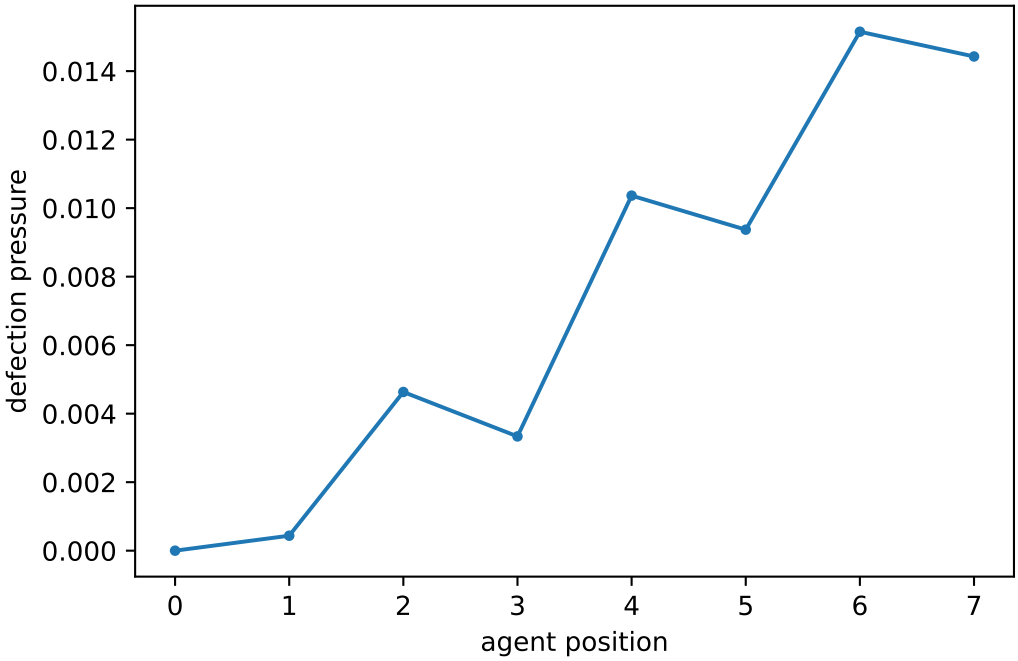

Figure 8.6: The difference between $P(A)_i$ for agent $i$ using $s_i$, minus her $P(A)_i$ using $s_{opt}$, for all $i$ positions in a group of 8 agents with $a = 0.4$ (simplified strategy). Higher positive values may indicate a stronger pressure to do so, while lower values a lower one.

In the previous section, it was demonstrated in Figure 8.4 that the presence of an uncooperative agent in a group of agents employing the aggregate strategy can potentially be harmful for the group as it may lead to a lower value of $\mathbb{E}(n_A)$; it was also clear that the presence of an uncooperative agent in such a group will not lead to a higher value of $\mathbb{E}(n_A)$ for the group. Moreover, the conditions (environmental uncertainty $a$ and position in the sequence $i$ of the uncooperative agent) that led to the highest benefit for her in terms of maximising her $P(A)_i$ were the ones that led the group to the lowest values of $\mathbb{E}(n_A)$.

This is not the case for groups of agents employing the majority strategy. As shown in Figure 8.7(b), there are several conditions (values of environmental uncertainty $a$ and position of the uncooperative agent $i$) for which the resulting $\mathbb{E}(n_A)$ of the groups is higher when an uncooperative agent is present, compared to when not; for example, this is the case for all values of $a$ when the uncooperative agent is is position $i = 1, 4, 5, 6, 7$ and not the case only for values of $a > 0.2$ when the uncooperative agent is in positions $i = 2, 3$. Moreover, as seen by comparing Figures 8.7(a) and 8.7(b), the conditions for which it is most beneficial for an agent to be uncooperative coincide with those that lead to the highest values of $\mathbb{E}(n_A)$ for a group with an uncooperative agent in that position.

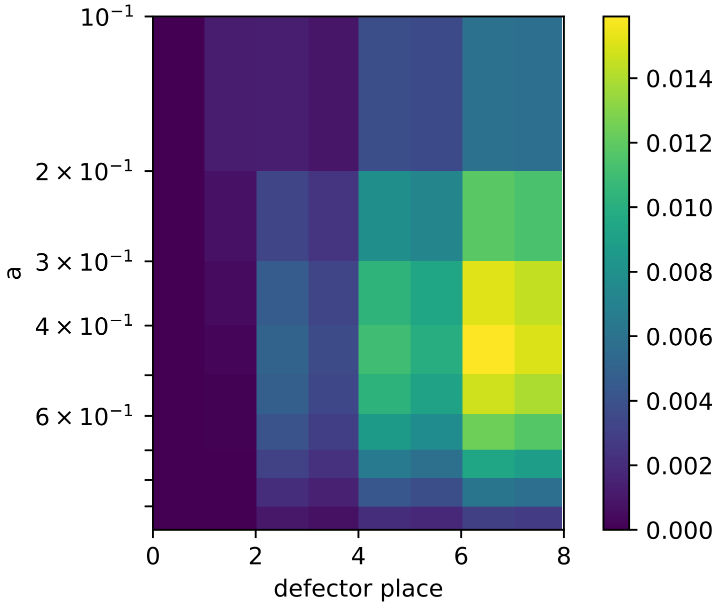

(a) comparing agents

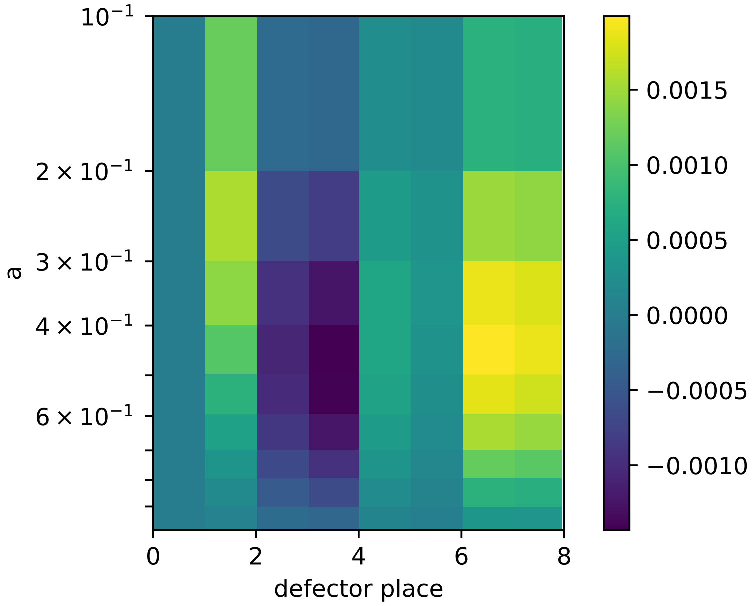

(b) comparing groups

Figure 8.7: Comparison of individual and group success when a selfish agent is present in position $i$ (simplified strategy). (a): difference between cooperative agent in place d, and defector in place d for d∈[0,8], in a range of $a \in(0,1)$. The plotted value is the cooperator's $p(A)_i$ minus the defector's $p(A)_i$ in the same position; zero means it makes no difference to defect, positive means it pays more to be selfish.
(b): difference between group with only cooperative agents, and group with defector in place d for d∈(0,8), in a range of $a \in(0,1)$. The plotted value is the clean group's $\mathbb{E}(n_A)$ minus the group with defector at place d $\mathbb{E}(n_A)$; zero means it makes no difference whether there is a defector, and negative means it's worse for the group to have a selfish agent in that position.

### Dynamic Strategy

Similarly to the aggregate strategy, the dynamic strategy also accumulates bias but as only the most recent agent is considered now, the bias is more discrete compared to the case of the aggregate strategy. This means that similar to the case of the majority strategy above, agents far along the sequence do not encounter over-amplified information, and can be more trusting of it compared to the case of the aggregate one. As a consequence, it is beneficial for agents to increasingly trust the available social information, as demonstrated in Figure 8.8, where agents evolve to increasingly larger values of $s_i$ as their position $i$ in the sequence increases.

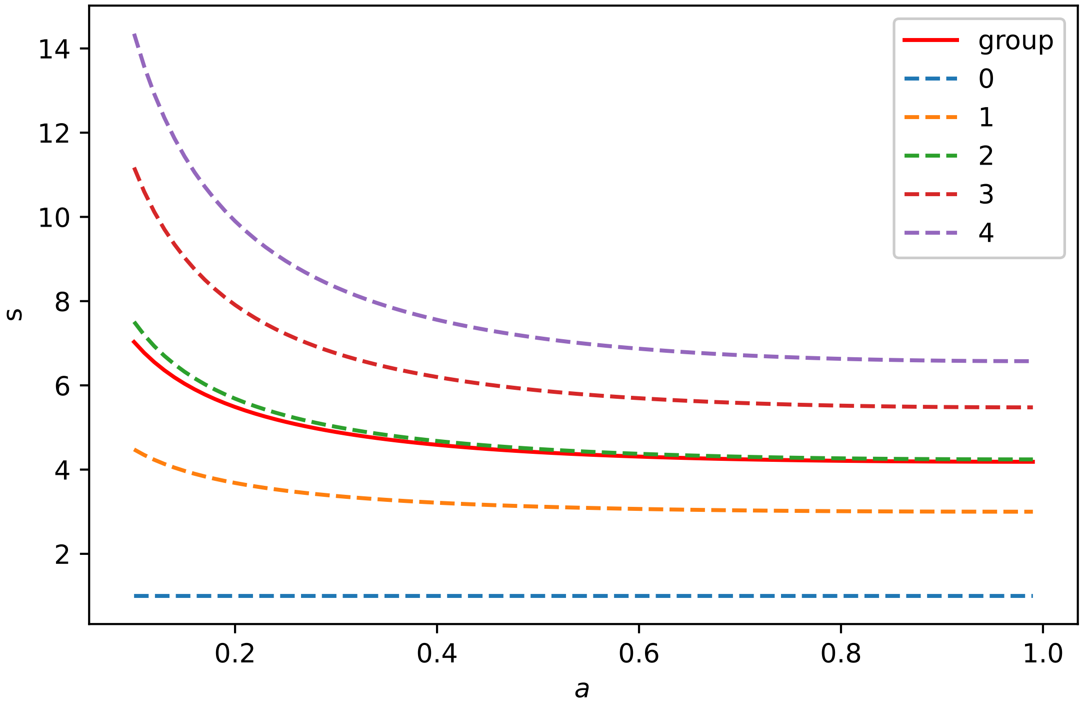

Figure 8.8: Collectively optimal $s$, and individually optimal $s_i$ for all agents in group of $n = 5$ using the dynamic strategy. The first agent (agent 0) has no available social information, hence her $s$ is irrelevant. For all the other agents, they all have values of $s_i$ larger than $s_{opt}$ with agent number 1 having the lowest one, and then increasing values of $s_i$ as rank increases.

Figure 8.8 demonstrates how the collectively optimal behaviour $s_{opt}$ is not personally optimal for any agent with fixed position within the group for the case of the dynamic strategy. In terms of the effect of position on $s_i$ for agents in places $i \geq 2$, the value of $s_i$ takes increasingly higher values. This is because the social information here is only provided by one agent (the most recent one), who is assumed to have assessed it enough herself, making it fairly trustworthy. Agent $i$=0 is the case of the agent without social information (for which the concept of sociality is nonsensical) and the case of the first agent (agent $i$=1) with social information, who can only observe a decision made entirely based of her conspecifics' personal information $a$, which she also shares: in her case, relying on social information doesn't provide significant additional insight for decision-making. But agent $i$=2 will now have access to two decisions, one of which social (hence providing updated information) leading to a higher value of $s_i$. For all other agents after her, the available social information is sufficiently weighted by the previous agent, making it reliable. So due to the nature of the dynamic decision-making strategy, for agents further back in the sequence the available social information is increasingly useful, leading to increasing values of $s_i$.

Despite the values of $s_i$ increasing with $i$, the same is not the case for the evolutionary pressure on the agents. The very first agent with access to social information has the highest evolutionary pressure to be uncooperative, as she has access to unbiased information from agent $i$=0. Agent $i$=2 has significantly lower evolutionary pressure, as now she is confronted with biased social information, and so is agent $i$=3: in both those cases, the previous agents are assumed to make very dependent decisions, making them not very informative. In the case of agent $i$=4 however, the evolutionary pressure is higher compared to the one on agent $i$=3 as now enough decisions have been made before her for the available social information to become more reliable. The reliability of social information increases for agents part this point, hence the evolutionary pressure on them increasing henceforth.

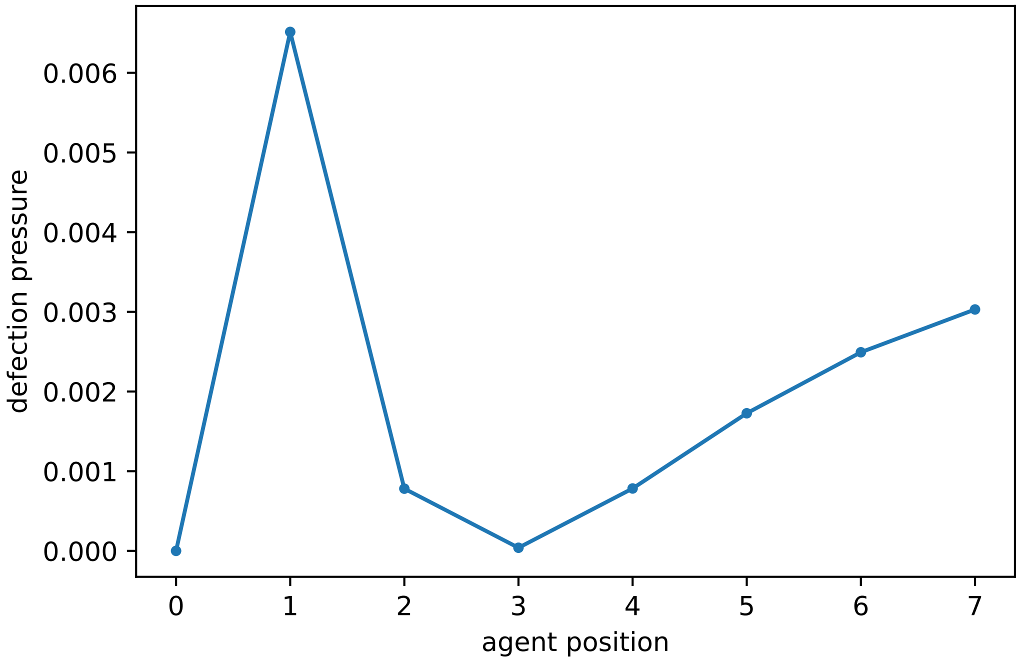

Figure 8.9: The difference between $P(A)_i$ for agent $i$ using $s_i$, minus her $P_{A_i}$ using $s_{opt}$, for all $i$ positions in a group of 8 agents with $a = 0.4$ (dynamic strategy). Higher positive values may indicate a stronger pressure to do so, while lower values a lower one.

In the two previous sections, it was demonstrated in Figures 8.4 and 8.7 that whether the presence of a uncooperative agent will be harmful to the group or not (in terms of lowering or increasing its collective measure $\mathbb{E}(n_A)$) depends on the strategy: in the case of the aggregate strategy, where the observed information by the focal decision-makers is the difference of prior decisions $\Delta n = n_A - n_B$ – in other words, when the available social information heavily propagates bias – the presence of an uncooperative agent has the potential to be harmful. In this case, when the agents comply to using $s_{opt}$, the result is a shared level of sociality that is fairly constrained in intensity (as discussed in depth in Chapter 5) which helps mitigate against propagation of mistakes; this results in individual values of $P(A)_i$ to also be fairly constrained. Hence when an individual agents employs a personally optimal sociality of $s_i$ to maximise her personal $P(A)_i$, she has a negative effect to this mitigating process. On the other hand, in the case of the majority strategy where the agents observe the sign of this difference ($sign(\Delta n)$) the impact of an uncooperative agent can be positive for many cases which can be attributed to the fact that agents using the majority strategy do not need to mitigate against bias propagation and as such the presence of an uncooperative agent is not as impactful.

In the case of the dynamic strategy, the presence of an uncooperative agent is not negative for any condition (value of environmental uncertainty $a$ and position of the uncooperative agent $i$) as shown in in Figure 8.10(b). The highest $\mathbb{E}(n_A)$ is observed when the uncooperative agent is in position $i$=1 – coinciding with the position with the highest evolutionary pressure for the agent. Directly following that, the benefit for the group in terms of $\mathbb{E}(n_A)$ is significantly lower when the uncooperative agent is in position $i$=2, and non-existent when she is in position $i$=3. From position $i$=4 onwards, the benefit for the groups is mostly positive (as opposed to none) and increasing. This coincides with the personal benefit of being uncooperative for the agent in that position.

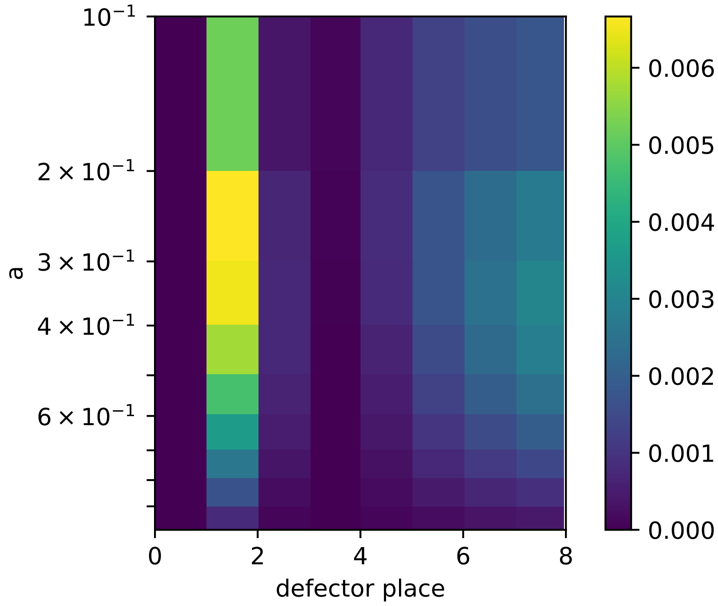

(a) comparing agents

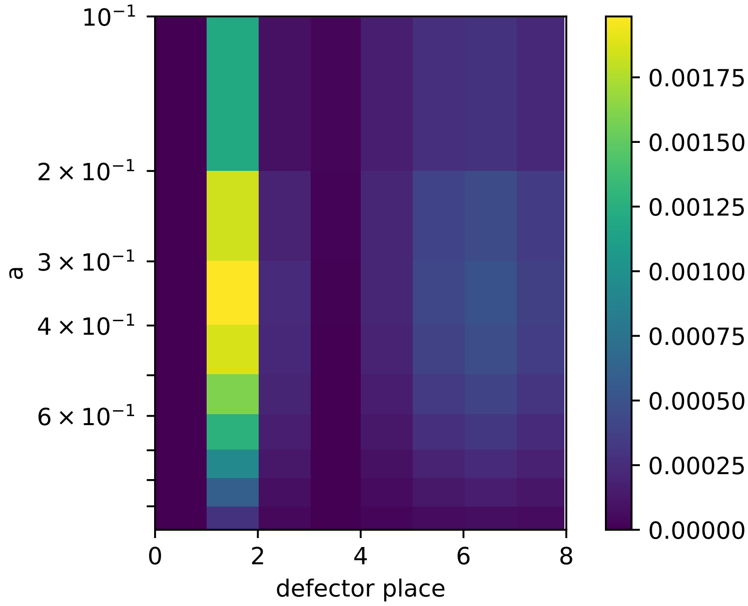

(b) comparing groups

Figure 8.10: (a): difference between cooperative agent in place d, and defector in place d for d∈[0,8], in a range of $a \in$[0,1]. The plotted value is the cooperator's $p(A)_i$ minus the defector's $p(A)_i$ in the same position: zero means it makes no difference to defect, positive means it pays more to be selfish.
(b): difference between group with only cooperative agents, and group with defector in place d for d∈[0,8], in a range of $a \in$(0,1). The plotted value is the clean group's $\mathbb{E}(n_A)$ minus the group with defector at place d $\mathbb{E}(n_A)$: zero means it makes no difference whether there is a defector, and negative means it's worse for the group to have a selfish agent in that position.

## 8.3 Individual Optimality in Evolutionary Stable Groups

Above we looked at what happens when one agent is optimising her $P(A)_i$ within a group that is collectively optimal (i.e. uses $s_{opt}$, which optimises $\mathbb{E}(n_A)$). We saw how there is always a distance between those two, yet this is to be expected, since as demonstrated in Sigalou and Mann, 2023 the collectively optimal group behaviour is never evolutionary stable. Here we repeat this analysis, to observe how the personally optimal behaviour of specific agents evolves, in groups that are evolutionary stable.

In evolutionary stable groups, we observe a similar pattern in the case of non-mixed groups: for a well-mixed group of 5 agents in an environment of $a = 0.4$ using the aggregate strategy, the evolutionary stable value of sociality is $s_{ESS} = 2.76$, and on average each agent $i$ has a probability of $P(A)_i = \mathbb{E}(n_A)$=0.7792 of choosing A over B (note that for $s_{opt}$ we have $P(A)_i = \mathbb{E}(A) = 0.78215 < 0.7792$, which is expected given that the evolutionary stable strategy is sub-optimal). For an unmixed group of 5 agents if they continue to employ $s = 2.76$ though, depending on their position the average probability that each agent $i$ will choose option A is going to be as follows:

| i | $P(A)_i$ |
|---|---|
| 0 | 0.71429 |
| 1 | 0.75967 |
| 2 | 0.79137 |
| 3 | 0.81031 |
| 4 | 0.82038 |

Table 8.3: Average probability of choosing option A, for agent in position i.

So now, instead of all agents expecting the same long-term outcome, each agent has a position-specific outcome – with some under-performing compared to the collective value of $\mathbb{E}(n_A)$, and some performing better. Note that 'average probability for agent i' is the arithmetic average of all the possible ways agent $i$ can choose A. Note that for the first agent with access to social information, $P(A)_1 = 0.75967 > 0.75622$, i.e. the $P(A)_1$ in the case of a collectively optimal group; for all following agents, their $P(A)_i$ values are lower in the case of an evolutionary stable group, compared to a collectively optimal group. Note that this is for a specific case of $a, s$ and strategy.

### 8.3.1 Aggregate

Below (Figure 8.11) are the plots for agents personally optimising in an evolutionary stable group, when both are using the aggregate strategy. All values of $s_i \neq s_{ESS}$ with some being lower and some higher that it, with the tendency being to lower the $s_i$ value as $i$ increases (with only the first few agents with social information being more social than is evolutionary stable). This is different compared to the case of a collectively optimal group (shown previously in Figure 8.1) where all values of $s_i$ were higher than $s_{opt}$.

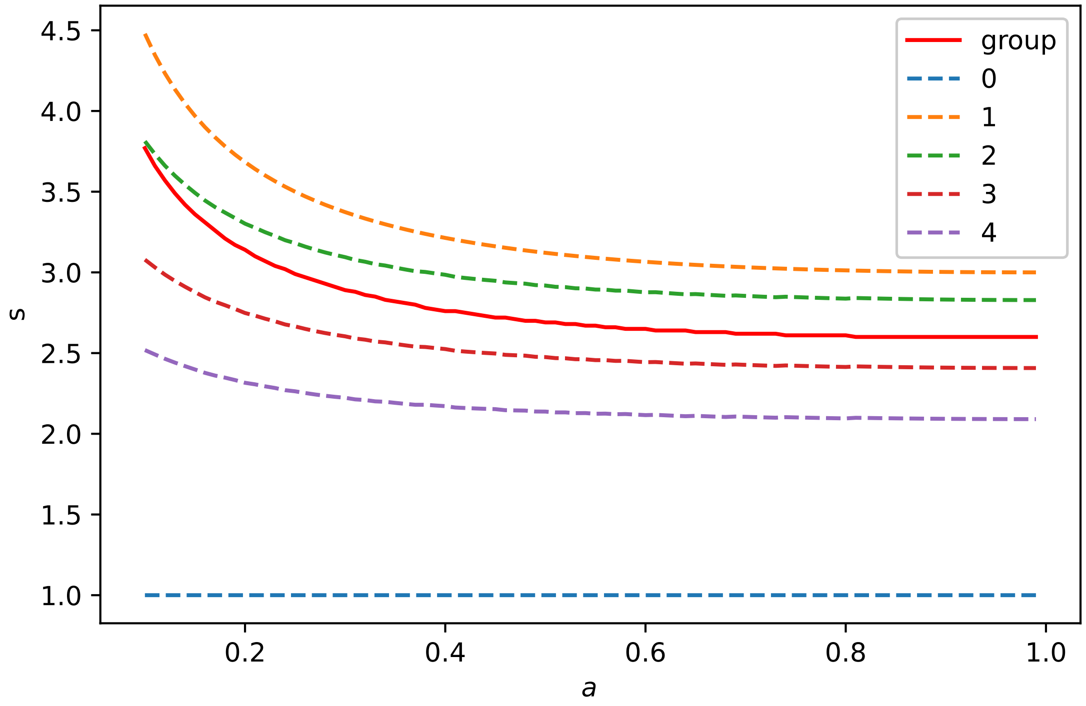

Figure 8.11: Evolutionary stable $s$, and individually optimal $s_i$ for all agents in group of $n = 5$ using the aggregate strategy and having evolutionary stable sociality. The first agent (agent 0) has no available social information, hence her $s$ is irrelevant and set as equal to 1. All subsequent agents have access to social information provided by agents that are optimising $\mathbb{E}(n_A)$.

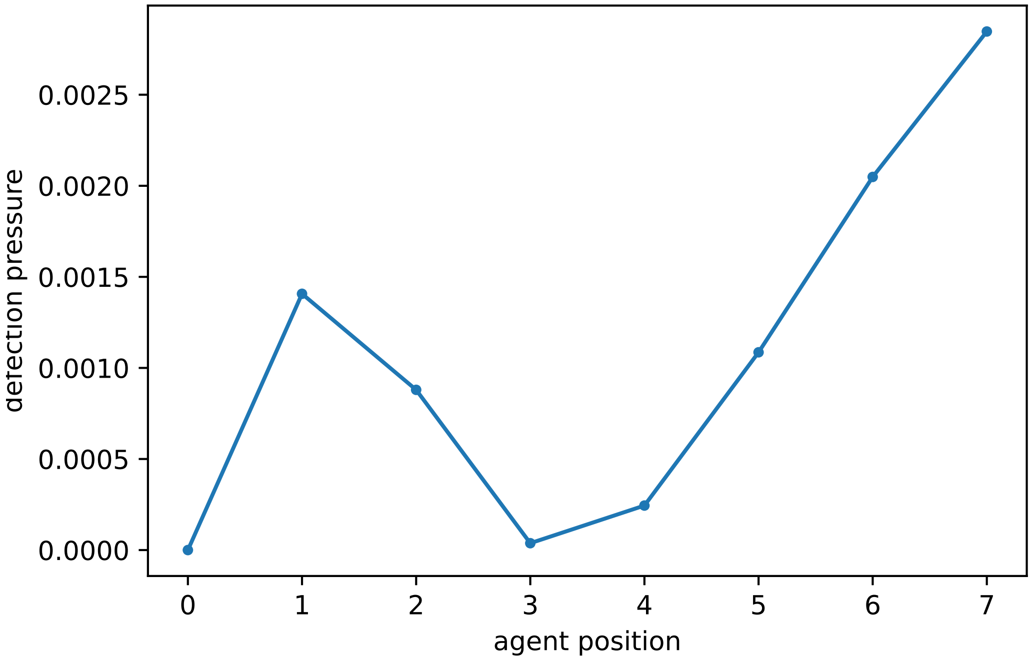

Figure 8.12: The difference between $P(A)_i$ for agent $i$ using $s_i$, minus her $P(A)_i$ using $s_{ESS}$, for all $i$ positions in a group of 8 agents with $a = 0.3$ (aggregate strategy). When said difference is positive, it's more profitable for agent i to use $s_i$ instead of $s_{ESS}$. Higher positive values may indicate a stronger pressure to do so, while lower values a lower one.

The evolutionary pressure also looks different, with an initial dip after the first agent with social information, and a subsequent increase after a few agents (Figure 8.12). This is very dissimilar to the evolutionary pressure for a group with same $a, s$ and size that is collectively optimal, where after a peak in evolutionary pressure for the first two agents with social information, the evolutionary pressure decreased as position in sequence $i$ increased, eventually reaching zero. This can be explained by considering each agent being deviating separately: agent $i$=1 has access to some unbiased social information from agent $i$=0, and hence makes sense that being more social pays off. Agent $i$=2 has access to the unbiased information from agent $i$=0 and the biased information from agent $i$=1, and hence while still paying off to be social she needs to be a little less to mitigate against potential bias amplification (which the aggregate strategy is prone to). Agent $i$=3 has access to more biased information compared to unbiased, and hence gets no benefit from being more social that the rest of the group –which is already over-social. Yet starting from agent $i$=4, there is an increase in evolutionary pressure. We are still in the aggregate strategy, ie prone to bias amplification, but now surprisingly it is beneficial to be even more over-social.

This is a general trend for all values of $a$, as shown in Figure 8.13(a). The phenomenon is less pronounced for extreme values of $a$ and more pronounced to medium-to-high, but overall it is more beneficial for agents further down the line to be non-cooperative and optimise personally rather than retain the group's behaviour. Unlike the case of the collectively optimal group (Figure 8.4), for the evolutionary stable group we notice that the positions $i$ where it is more favourable for an agent to be non-cooperative as not the positions where a deviator would negatively impact the group; rather, the most impactful position for a defector to be in is $i = 1$ which is a position a fairly constrained benefit compared to future ones. On the other hand, the position that is more beneficial for the individual agent has a low positive effect for the group. The most beneficial defector positions for the group have fairly low evolutionary pressure.

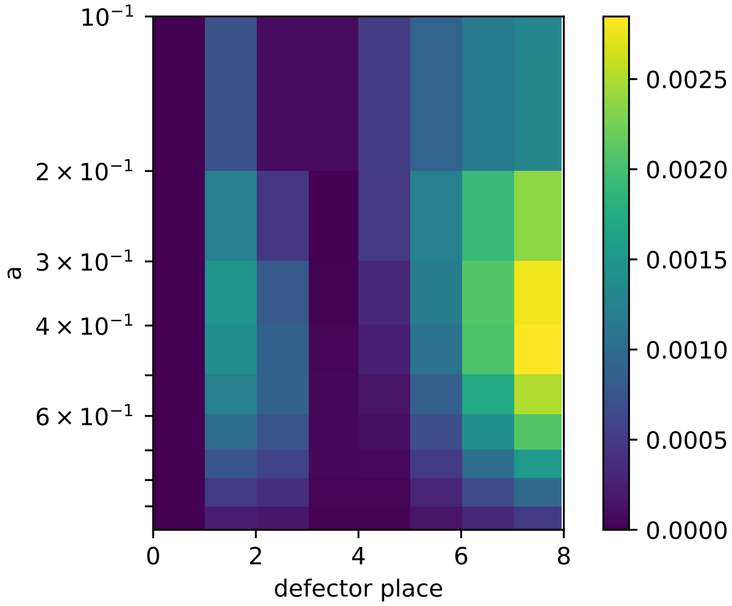

(a) comparing agents

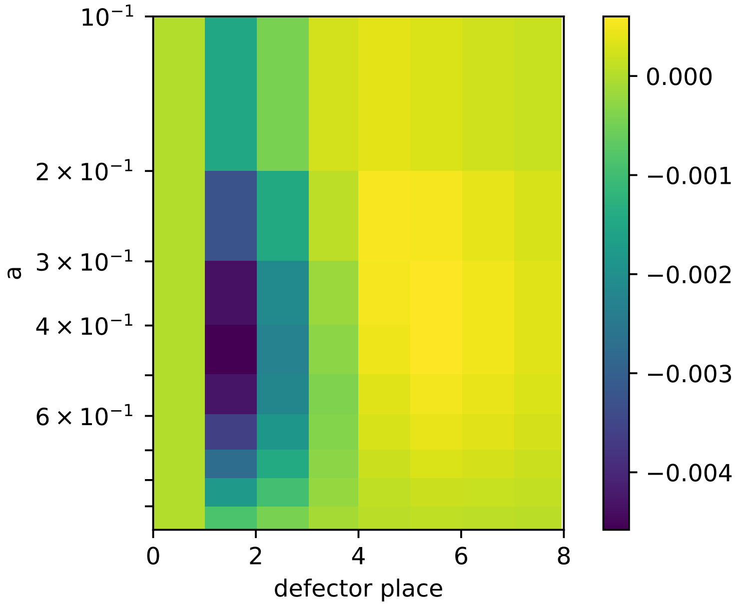

(b) comparing groups

Figure 8.13: Comparison of individual and group success when an uncooperative agent is present in position $i$ (aggregate strategy). (a): difference between uncooperative agent in place d, and cooperative agent in place d for d∈[0,8], in a range of $a \in(0,1)$. The plotted value is the uncooperative agent's $p(A)_i$ minus the cooperative agent's $p(A)_i$ in the same position; zero means it makes no difference to be uncooperative, positive means it pays more to be so.
(b): difference between group with an uncooperative agent in place d for d∈[0,8] and group with only cooperative agents, for a range of $a \in(0,1)$. The plotted value is the $\mathbb{E}(n_A)$ of the group with an uncooperative agent in position $i$, minus the $\mathbb{E}(n_A)$ of the group without an uncooperative agent in that position; zero means it makes no difference whether there is an uncooperative agent, and negative means it's worse for the group to have an uncooperative agent in that position.

Overall, by comparing Figures 8.13(a) and 8.13(b) we notice that it is most likely that a defector will have a negligible or positive impact for the group: the highest evolutionary pressure is for agent $i$=8 at a value of $0.4 < a < 0.6$, which will not affect the value of $\mathbb{E}(n_A)$; the most impactful case for the group would be that of agent $i$=1 'defecting', which is a case with quite low evolutionary pressure.

### 8.3.2 Majority

Below are the plots for agents personally optimising in an evolutionary stable group, when both are using the majority strategy. All values of $s_i \neq s_{ESS}$ with some being lower and some higher that it compared to the case of a collectively optimal group where all values of $s_i$ were higher (Figure 8.5), with the tendency being to increase the $s_i$ value as $i$ increases (but not in a 'linear' as, as a pattern similar to that of the collectively optimal group occurring.

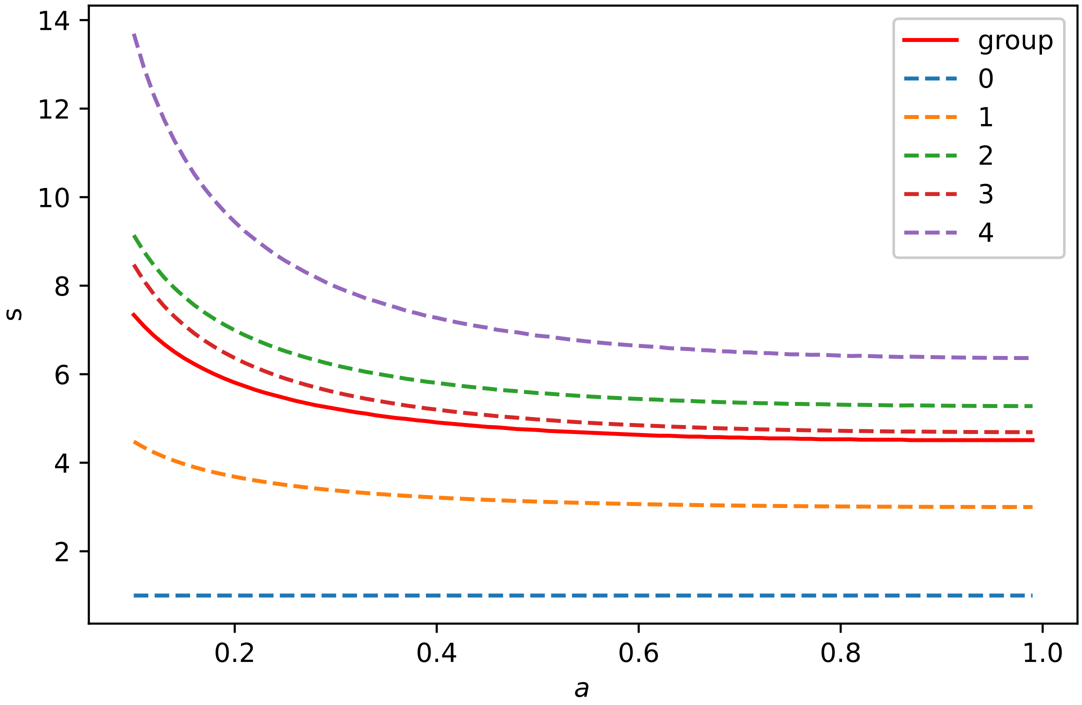

Figure 8.14: Evolutionary stable $s$, and individually optimal $s_i$ for all agents in group of $n = 5$ using the simplified strategy. The first agent (agent 0) has no available social information, hence her $s$ is irrelevant. For all the other agents, they all have values of $s_i$ larger than $s_{ESS}$.

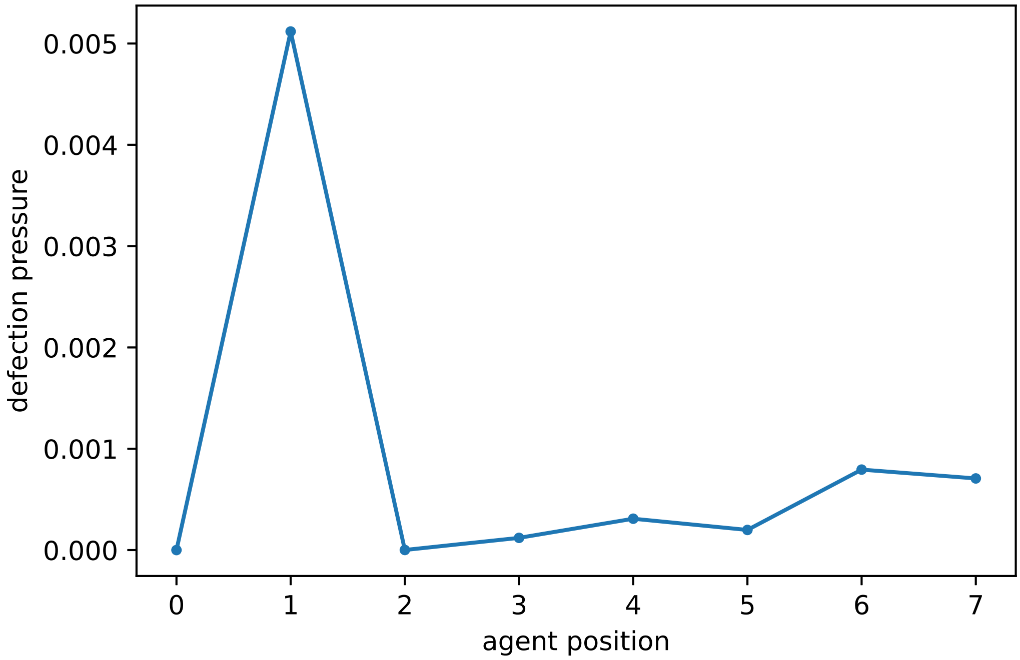

Figure 8.15: The difference between $P(A)_i$ for agent $i$ using $s_i$, minus her $P(A)_i$ using $s_{ESS}$, for all $i$ positions in a group of 8 agents with $a = 0.4$ (simplified strategy). Higher positive values may indicate a stronger pressure to do so, while lower values a lower one.

Again, in the case of the evolutionary stable group the evolutionary pressure also looks different to the one from the case of the collectively optimal group (Figure 8.15). Now the first agent with social information has the highest evolutionary pressure out of all agents in the group; this is followed by no evolutionary pressure for the second agent with social information, and following that all subsequent agents have a slightly increasing evolutionary pressure.

This makes sense if we consider each agent being the deviating one at a time: agent $i$=1 has access to some unbiased social information from agent $i$=0, and hence makes sense that being more social pays off. Agent $i$=2 has access to the unbiased information from agent $i$=0 and the biased information from agent $i$=1; in this case being more social than the group does not contribute a higher benefit, as the decision of agent $i$=1 is not informative enough. Agent $i$=3 has access to the unbiased information from agents $i$=0 and $i$=2, and hence gets some benefit from being more social that the rest of the group. Since the majority strategy does not amplify bias down the line, from this point forward it makes sense for each agent to be increasingly more social (since all agents from $i$=3 onwards have enough social information without bias amplification). This is a general trend for all values of $a$, as shown in Figure 8.16(a).

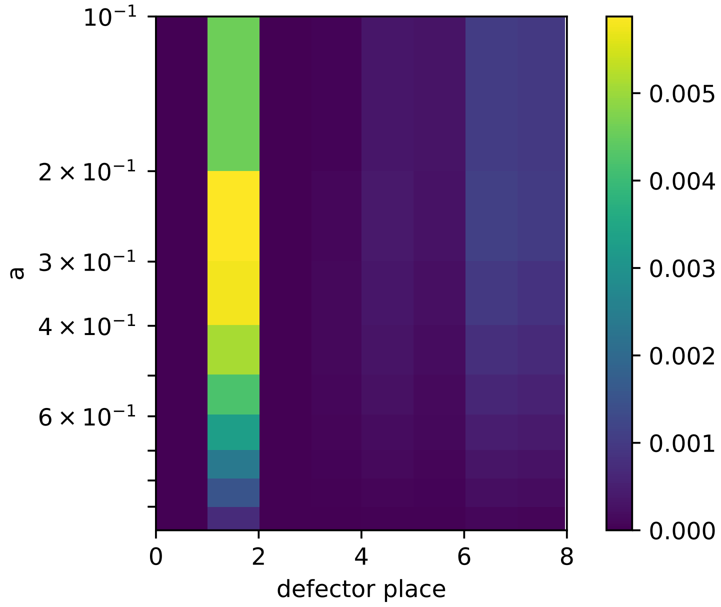

(a) comparing agents

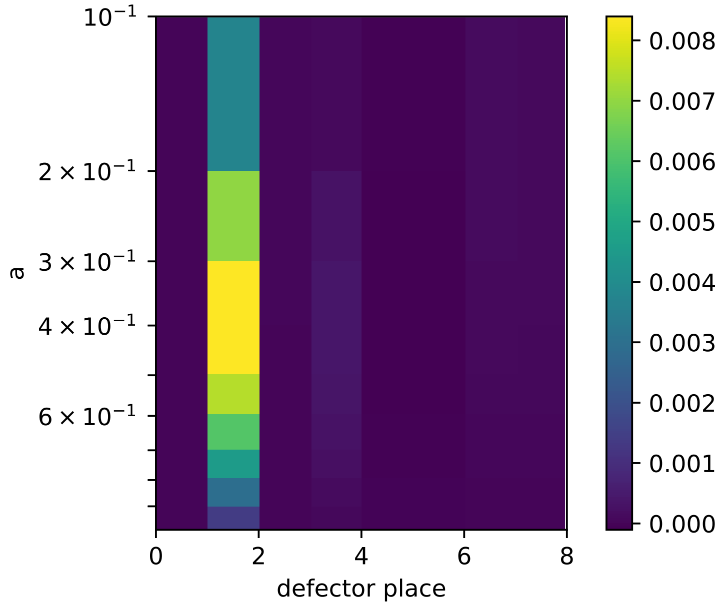

(b) comparing groups

Figure 8.16: Comparison of individual and group success when a selfish agent is present in position $i$ (simplified strategy). (a): difference between cooperative agent in place d, and defector in place d for d∈[0,8], in a range of $a \in (0,1)$. The plotted value is the cooperator's $p(A)_i$ minus the defector's $p(A)_i$ in the same position; zero means it makes no difference to defect, positive means it pays more to be selfish.
(b): difference between group with only cooperative agents, and group with defector in place d for d∈(0,8), in a range of $a \in (0,1)$. The plotted value is the clean group's $\mathbb{E}(n_A)$ minus the group with defector at place d $\mathbb{E}(n_A)$; zero means it makes no difference whether there is a defector, and negative means it's worse for the group to have a selfish agent in that position.

Overall, by comparing Figures 8.16(a) and 8.16(b) we notice that agent $i$=1 has the highest evolutionary pressure to personally optimise, which corresponds to a case that will be the most beneficial for the group overall as it translates to a higher value of $\mathbb{E}(n_A)$. All other agents have a negligible evolutionary pressure, and the effect of the group would be equally negligible or non-existent.

### 8.3.3 Dynamic

Below are the plots for agents personally optimising in an evolutionary stable group, when both are using the dynamic strategy. All values of $s_i \neq s_{ESS}$ with some values of $s_i$ being lower and some higher than $s_{ESS}$ (Figure 8.17), with the tendency being to increase the $s_i$ value as $i$ increases (in a pattern similar to that of the collectively optimal group –see Figure 8.8).

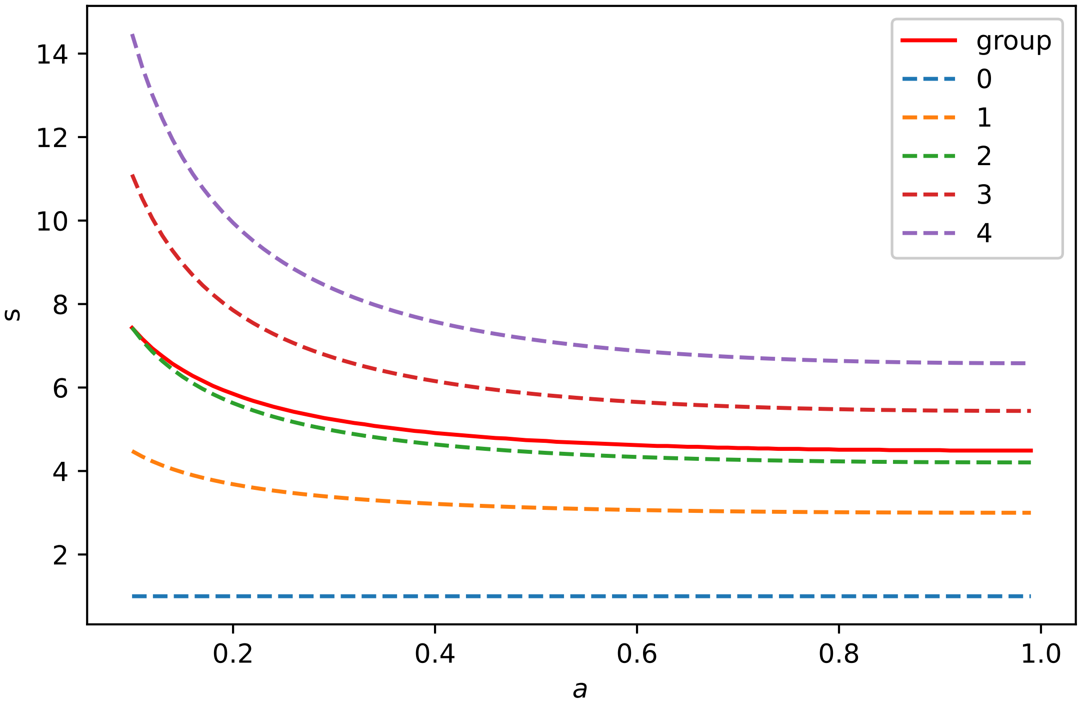

Figure 8.17: Evolutionary stable $s$, and individually optimal $s_i$ for all agents in group of $n = 5$ using the dynamic strategy. The first agent (agent 0) has no available social information, hence her $s$ is irrelevant. For all the other agents, they all have values of $s_i$ larger than $s_{ESS}$ with agent number 1 having the lowest one, and then increasing values of $s_i$ as rank increases.

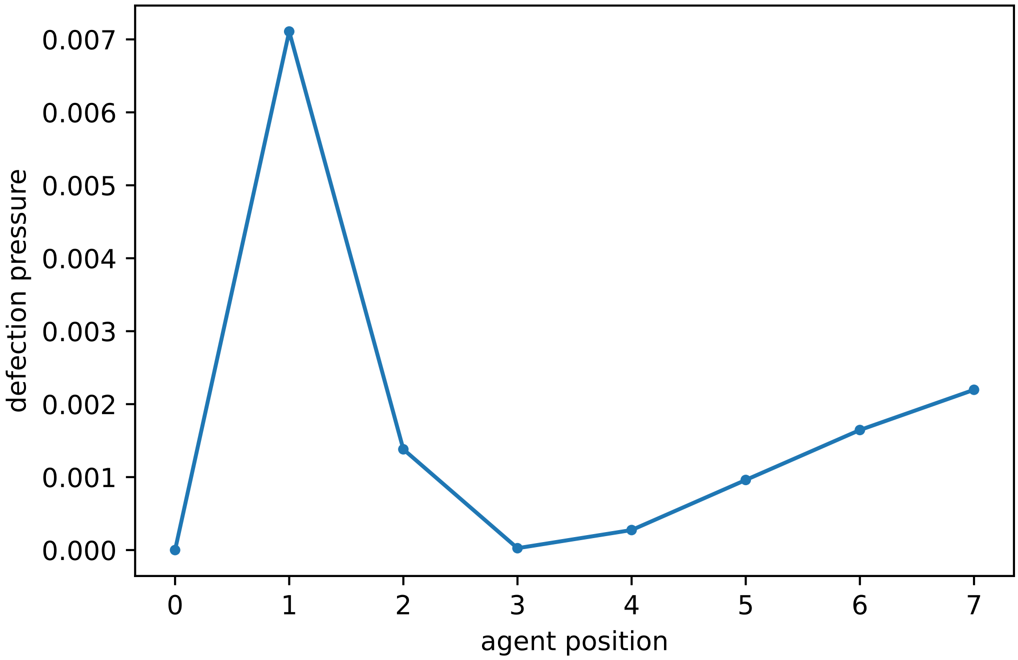

Figure 8.18: The difference between $P(A)_i$ for agent $i$ using $s_i$, minus her $P_{A_i}$ using $s_{ESS}$, for all $i$ positions in a group of 8 agents with $a = 0.4$ (dynamic strategy). Higher positive values may indicate a stronger pressure to do so, while lower values a lower one.

Again, in the case of the evolutionary stable group the evolutionary pressure also looks different to the one from the case of the collectively optimal group (Figure 8.18). Now the first agent with social information has the highest evolutionary pressure out of all agents in the group; this is followed by no evolutionary pressure for the second agent with social information, and following that all subsequent agents have a slightly increasing evolutionary pressure.

This makes sense if we consider each agent being the deviating one at a time: agent $i$=1 has access to some unbiased social information from agent $i$=0, and hence makes sense that being more social pays off. Agent $i$=2 has access to the unbiased information from agent $i$=0 and the biased information from agent $i$=1; since this is the case of the dynamic strategy, the assumption is that agent $i$=1's choice incorporates the information from agent $i$=0 and hence is more informative from the information she would be providing if using one of the static strategies (ie the aggregate or the majority) and so it makes sense for agent $i$=2 to be more social than the group, but less social than agent $i$=1 since the social information she has access to is more biased. Agent $i$=3 has now access to more biased than unbiased information and hence gets no benefit from being more social that the rest of the group. From that point onwards (ie from agent $i$=4 on), there is enough trustworthy social information (which is not amplifying bias, as this is not a property of the dynamic strategy) and hence there is an increased benefit to being more social than the rest of the group. This is a general trend for all values of $a$, as shown in Figure 8.19(a).

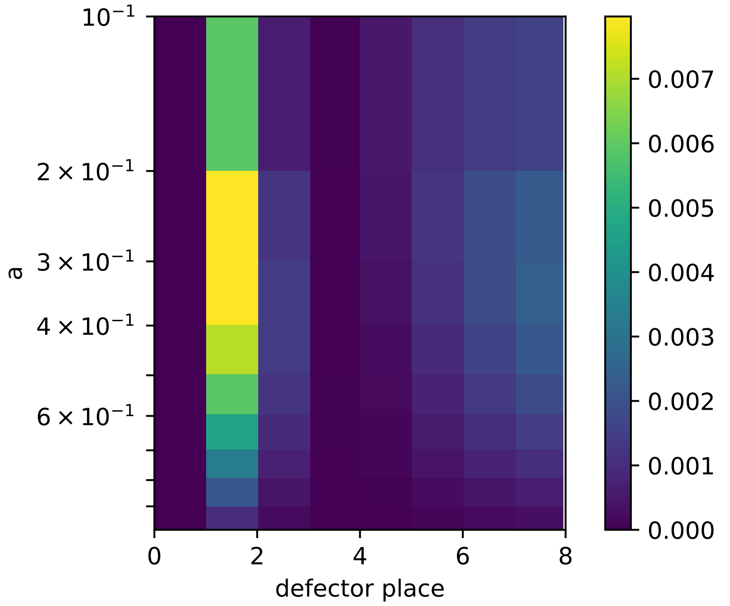

(a) comparing agents

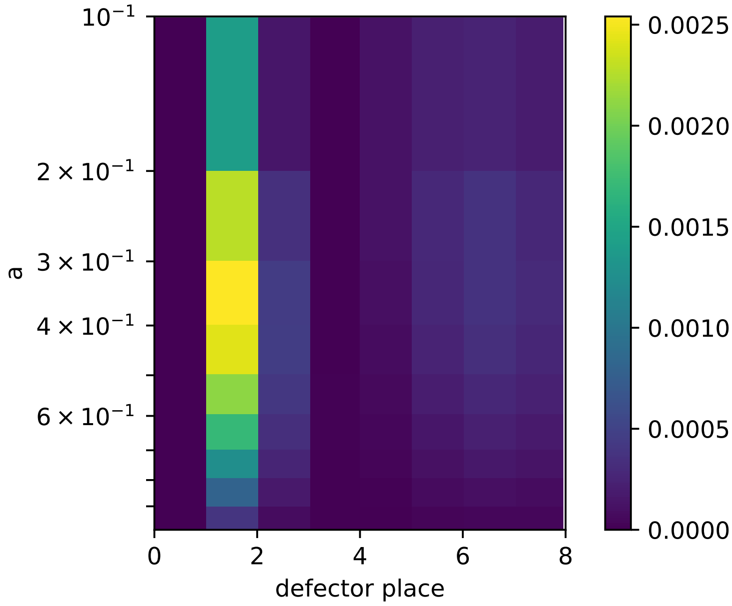

(b) comparing groups

Figure 8.19: (a): difference between cooperative agent in place d, and defector in place d for d∈[0,8], in a range of $a \in[0,1]$. The plotted value is the cooperator's $p(A)_i$ minus the defector's $p(A)_i$ in the same position; zero means it makes no difference to defect, positive means it pays more to be selfish.
(b): difference between group with only cooperative agents, and group with defector in place d for d∈[0,8], in a range of $a \in(0,1)$. The plotted value is the clean group's $\mathbb{E}(n_A)$ minus the group with defector at place d $\mathbb{E}(n_A)$; zero means it makes no difference whether there is a defector, and negative means it's worse for the group to have a selfish agent in that position.

Overall, by comparing Figures 8.19(a) and 8.19(b) we notice that agent $i$=1 has the highest evolutionary pressure to personally optimise, which corresponds to a case that will be the most beneficial for the group overall as it translates to a higher value of $\mathbb{E}(n_A)$. All other agents have a negligible evolutionary pressure, and the effect of the group would be equally negligible or non-existent.

## 8.4 Effect of Group Size

The collectively optimal and evolutionary stable values of $s$ depend on group size, with a tendency to decrease as group size increases. This is due to the fact that they are calculated as the best strategy on average for an agent assuming all possible positions within a group: hence for larger groups, bias propagation must be taken into account since it can negatively impact later positions, thus driving these values down (Figure 8.20).

(a) $s_{opt}$, $a = 0.3$ (b) $s_{opt}$, $a = 0.7$

(c) $s_{ESS}$, $a = 0.3$ (d) $s_{ESS}$, $a = 0.7$

Figure 8.20: Dependence of $s_{opt}$ and $s_{ESS}$ on group size for the aggregate strategy.

The majority strategy on the other hand don't face the same issue with bias amplification. For that reason the values of both $s_{opt}$ and $s_{ESS}$ is higher compared to the corresponding values for the aggregate one, with $s_{ESS}$ additionally having an increasing (rather than decreasing) trend (Figure 8.21).

(a) $s_{opt}$, $a = 0.3$

(b) $s_{opt}$, $a = 0.7$

(c) $s_{ESS}$, $a = 0.3$

(d) $s_{ESS}$, $a = 0.7$

Figure 8.21: Dependence of $s_{opt}$ and $s_{ESS}$ on group size for the majority strategy.

Either way, both for the aggregate and the majority strategies it appears that the values of $s_{opt}$ and $s_{ESS}$ 'slow down' as group size increases, possibly indicating a plateau. Since these results are preliminary, this is simply a direction for future research rather and should not –at this point– be taken at face value.

## 8.5 Discussion

This chapter explores the evolution of sociality of individual agents within a group, when assuming unmixed groups. In unmixed groups, each agent is fixed in one position, therefore it no longer makes sense to consider the average behaviour over time as now the agents' performance during decision-making is linked to their position $i$ inside the sequence.

I assume groups with agents in fixed positions, where every agent but one (whom I call uncooperative) is employing the collectively optimal level of sociality $s_{opt}$. In that case the uncooperative agent, by deviating from the commonly used $s_{opt}$ is able to achieve a higher value of $P(A)_i$ compared to when abiding to $s_{opt}$. I found that regardless of the uncooperative agent's position $i$ in the sequence, and regardless of the strategy employed by the group, the uncooperative agent benefits most by employing a personally optimal value $s_i$ that is different from $s_{opt}$. This is to be expected, since the collectively optimal value of $s$ is not evolutionary stable and as such vulnerable to deviators (shown in more detail in Chapter 5). While it is clear that $s_i \neq s_{col}i$, the exact way is not linear; there is a variety depending on the strategy (for instance, the trend is quite different for the aggregate and the majority strategy), while specifically for the aggregate strategy there is an overlap for $s_i$ for very low values of $a$ (see Figure 8.1). Further work is require to provide a deeper understanding of these.

A similar behaviour is observed when considering agents personally optimising within an evolutionary stable group: in this case, agents adopting $s_i \neq s_{ESS}$ are able to achieve a higher value of $P(A)_i$ for all positions and strategies considered. Unlike in the case of the collectively optimal group, this is not as obvious as not the group's value of $s = s_{ESS}$ in evolutionary stable, meaning that there is no value of $s \neq s_{ESS}$ able to outperform $s_{ESS}$. For this reason, we expect that the uncooperative agents do not remain in these for long. In other words, imagine $s_6$ being the most successful strategy in a group with originally $s_{ESS} \neq s_6$ and $s_6$ eventually becoming dominant. I expect that such a group would then be susceptible to invasion by other $s' \neq s_6$ and will eventually return to the original group's $s_{ESS}$.

The effect of an uncooperative agent in position $i$ was also explored. Agents early on in the sequence have the potential to be more impactful for the overall group performance as their decisions affect a large proportion of the group. The impact an uncooperative agent has on the group's resulting collective measure $\mathbb{E}(n_A)$ depends on that agent's position $i$, and whether the group is at a collectively optimal or evolutionary stable condition. For the case of a collectively optimal group, when the decision-making strategy allows for bias propagation then the effect of an uncooperative agent can either be harmful or neutral for the overall group (depending on her place in the sequence), but if not then her effect can lead to a higher $\mathbb{E}(n_A)$. For the case of an evolutionary stable group, in the case of the aggregate strategy it is possible for the uncooperative agent's effect to be negative for the group, but in a different way compared to a collectively optimal group as now the evolutionary pressure for each position is vastly different; for the case of the majority and the dynamic strategy however, the impact of an uncooperative agents can either be beneficial of neutral.

To a large degree, whether the presence of an uncooperative agent will be negative towards $\mathbb{E}(n_A)$ or not – as well as by how much – depends on the quality (reliability) of social information she is transmitting. When the decision-making strategy is prone to over-amplifying bias, then it is mostly negative: in this case, the relationship between group and individual agents takes the form of a competition, additionally so because of the coincidence between the position with the highest evolutionary pressure being the most harmful one for the group. On the other hand, when the employed strategy does not lead to an intense propagation of bias the presence of an uncooperative agent can often times have a positive effect for the rest of the group. In these cases there is no longer a conflict of interest between agent and group; in some cases, it can even be equally beneficial for the agent in the position under the most evolutionary pressure to employ $s_i$, and to the group she is in.

A clear effect of environmental uncertainty in evolutionary pressure & effect of uncooperative agent is also demonstrated in the above analysis. Environments of low uncertainty (low values of $a$) mean that regardless of place in the sequence or employed strategy, the agents have more reliable personal information; this also means that the available social information will be equally more reliable compared to cases of higher values of $a$, as it results in prior decisions being more accurate. This results in generally lower evolutionary pressure to be uncooperative for lower values of environmental uncertainty compared to higher values, as in this case they already obtain high enough values of $P_{A_i}$ and an increase in sociality wouldn't increase it as mush as it would for more uncertain environments. To this there is one exception, that of the first agent with access to social information employing the aggregate strategy, due to the nature of this strategy. On the other hand environments with high uncertainty (high values of $a$) result in poor personal information, and subsequently less reliable social information, as now the previous decisions are more prone to mistakes. This results in lower evolutionary pressure compared to lower values of $a$, as in that case the agents only have access to very unreliable social information, to which there is no benefit in over-trusting as that would result in poorer values of $P_{A_i}$.

This chapter provides some insight into the evolutionary dynamics of individual agents in groups, as well as into the relationship between agent and group and its dependence on several factors. It is worth to note at this point, two elements missing from this analysis that are the logical next steps for future research. The effect of group size has not been explored enough here. Group size is an important factor in social groups; in this case it is especially important as the group size $N$ determines what $S_{opt}$ is, so it also affects the values $s_i$, and as shown in the preliminary results there might be a limit for its value as group size increases. Different group consistencies are also not considered here, and notably groups where everyone has evolved to be personally optimal.

For unmixed groups, a reasonable assumption is that the agents will individually adapt to their specific positions. So in a group where everyone has evolved to be uncooperative (i.e. optimise their $P_{A_i}$ instead of $\mathbb{E}(n_A)$), the focal decision-maker $j$ has a different problem: she now has access to increasingly dependent social information. Since the agents now care solely about the information they consume (and not about the information they provide) they follow more strongly, hence not utilising their own personal information. So agents further back in line will no longer have access to an added layer of information, like in the previous case. The expectation is that in this case, the agents will evolve to a different value of personally optimal behaviour $s'_i \neq s_i$.

# Chapter 9

# Conclusions

The work presented in the chapters above (and especially in Chapters 5, 7, 6 and 8 that present the original work conducted during this PhD) used a binary decision-making model to explore some fundamental aspects of social decision-making. This final chapter summarises the main contributions of this work, and outlines some directions for future work that naturally follow from it.

## 9.1 Main Contributions

### 9.1.1 Evolutionary Stability

Rather than viewing sociality (in this case, following social information during decision-making) as just an observed trait, this thesis considers it an evolved one in the principle that 'nothing in biology makes sense except in the light of evolution' [Dobzhansky, 2013] and as such observed behaviours of animals (and groups of animals) should always been taken as *evolutionarily rational* instead of plainly rational. This approach allows for the study of the evolution of sociality in groups navigating uncertain environments, with its main contribution being the analysis of the long-term effects this has on groups obeying an intuitive adaptation process.

This added attribute is a simple yet important one. While sequential binary decision-making has a variety of appropriate applications, some considering only short time-scales, both it and the chosen model have been used in contexts where this adaptation process is important. By adding this assumption, I was able to determine a necessary constraint on the level of sociality of agents tied to the environmental uncertainty, which was not previously considered. By doing so, I have provided an important guideline for experimental researchers attempting to apply their data to this model.

Moreover, by considering a purely theoretical model I was able to gain generalisable insight into adaptive behaviours that have previously been characterised as 'irrational', such as the tendency of animal groups to become over-social (instead of evolving to be optimal). By generating the long-term evolutionary tendencies of groups without predisposing them to behave in an expected way, I was able to contextualise this over-sociality as an adaptive behaviour rather than a mystery.

### 9.1.2 Use of Social Information

An important theme throughout this work is the use of social information provided by conspecifics, during sequential binary decisions. This work analysed the effect that different decision-making rules have on the evolution of sociality of heterogeneous groups, as well as the role of group size in the existence of an evolutionary stable strategy among the ones considered here. Lastly it observes the existence of such a strategy in the case of very large groups, where the effect of small perturbations (in the form of a single agent employing a different strategy) is less significant for the overall group. Additionally Chapters 5 and 6 explored the effect of different decision-making strategies in the final probabilities of different configurations between the available options; it was shown computationally in Chapter 5 and analytically in Chapter 6 that strategies that allow for the accumulation of bias over time (and hence for the emergence of consensus decisions, such as all agents of the group choosing the same option) are also the ones that are prone to cascades.

The final contribution of this work is in regards to the relationship between agents in a group. By considering unmixed groups where the agents remained in fixed positions within the sequence it was possible to consider the variance in personal sociality between the different positions; since each position has access to more (or fewer) prior decisions, this directly translates to the reliability of this social information. Due to this effect, the evolutionary pressure for agents to personally optimise their sociality instead of the one of the group (as it is the case in well-mixed groups) differs between positions within the group, and so does the effect of a deviating agent for all the following decisions and the resulting collective group outcome. By exploring this effect for the three decision-making rules considered throughout this work (the aggregate, the majority and the dynamic) I showed that whether the deviating agent compromises the rest of the group (i.e. whether her impact is negative for the resulting collective outcome compared to the case where she employs the collectively employed level of sociality) depends on the decision-making strategy. Specifically I found that this is only the case for the aggregate rule, where social bias has an important and increasing effect, while in the case of the other two rules the effect of a deviating agent can even have a positive effect on the group. This result is important as it challenges the notion of direct competition between group and individual, as well as the notion that there is an inherent conflict in goals between the two. This result, stemming from an evolutionary analysis with very minimal initial conditions provides a wider framework to understand these dynamics.

## 9.2 Directions for Future Work

In this work I assume groups of identical agents. These agents navigate an environment with uncertainty $a$ and have an adaptable trait described by the parameter $s$. The evolution of trait $s$ is driven by the fitness of the decision-making rule being used: in this context, fitness refers to the value of parameter $s$ that results in the highest value of choosing the best option. This value is calculated computationally using python's 'optimize' function. The concept of a rare mutant taking over the population is applied following the above. The mutant will have a relatively better fitness if her value of $s$ grants a higher probability for choosing the best option, compared to if she employs the strategy of the group. Starting from this model, we can formulate several naturally occurring questions, here separated by theme.

### 9.2.1 Group Composition

Only homogeneous groups were considered here, yet many groups in nature are heterogeneous with many operating on a distinction of 'leaders' and 'followers', where the former is investing in acquiring information and making decisions, and the latter in social behaviour (i.e. following the social cues provided by the leaders) [Guttal and Couzin, 2010]. Research done on the matter usually work with theoretical models build on observations or actual data, and aim at understanding the effect of heterogeneity in the groups, and how this contributes to these groups demonstrating collective behaviour [Aplin et al., 2014]. The framework I propose here can contribute to our understanding of heterogeneity in groups: one direction is to create groups where agents have different personal information $a$ and investigate how this affects the collectively optimal and evolutionary stable values of sociality ($s_{opt}$ and $s_{ES}$ respectively) in the case of a well mixed group, as well as the personally optimal $s_i$ for unmixed groups.

Additionally, this framework can be used to investigate the emergence and maintenance of heterogeneous groups. Heterogeneity can lead to weaker social cohesion and eventual group fission [Conradt and Roper, 2000], while sourcing information from agents with different preferences can be non-informative [Mann, 2020]; nevertheless heterogeneous groups are observed in nature. Using the approach outlined in this thesis, one can explore heterogeneity as the result of balancing personal and collective optimality, for instance by finding a case where the agents of a group increase the collective measure $\mathbb{E}(n_A)$ as a result of optimising their personal $P(A)_i$.

### 9.2.2 Cost

Additionally I only considered cases where there was no cost in employing different decision-making rules, yet in nature there is usually a trade-off between more or better information and cost (in the form of an increase in the required processing time, or a larger body size and caloric needs). Extensions to this model may additional investigate reasonable cost functions as an added feature to the decision-making strategies and study the differences in outcome between both the analysis of the distinct strategies, as well as the dynamics between them. For example, it is reasonable to expect that the aggregate strategy – which has so far been the evolutionary stable one for infinite sized groups – may cease to be it as intuitively it should be more costly to employ compared to both the majority and the dynamic. We expect a difference in the dynamics between the strategies in the case of finite groups as well; assuming that the costs are proportional to complexity, there may be a clearer pattern in terms of which other strategies can invade each.

### 9.2.3 Evolutionary Dynamics

This work was primarily computational and based on a specific model, however many of the concepts used for the assumptions stem from a different framework: that of evolutionary game theory. Evolutionary game theory provides a blueprint for studying the evolution of interactions in groups, where they are framed as pairwise interactions and analysed using either differential (in the case of infinite groups) or stochastic equations (in the case of finite groups).

By abiding to the formal formulations of Evolutionary Game Theory, more generalisable analytical results can be obtained regarding the main themes of this study: invasion and evolutionary stability. For instance, Chapters 5 and 7 examine the invasion of a single invader in an otherwise homogeneous groups and found the existence of an equilibrium for this process. A more complete analytical approach, such as using a fitness equation such as Equation 3.12 the existence of more equilibrium points can be explored between two strategies in the case of an infinite size group.

The computational nature of the present analysis also left some blind spots such as the amount of 'time' needed for said equilibrium to be reached. This can be complemented if the analysis is instead performed by using a Moran process (presented in a fair amount of detail in Chapter 3). In that way, by determining different payoffs and simulating the games we can obtain results that will additionally include the amount of time needed to reach the equilibria. Using a Moran process additionally allows to consider more complex interactions. Costs, which are absent from the model analysed in this thesis, can be included an the effect they have on the evolutionary stability of the different decision rules can be extracted and compared to the results of the present work. Variations of neutral and random drift can also be considered to phrase strategies with different fitness, and a more robust comparison can be made between them.

The influence of group size can also be explored in a more consistent way. Evolutionary game theory provides separate frameworks for working with infinite and finite groups, based on the principle that in the latter stochasticity has a significant effect which in the case of infinite groups is negligible and may be ignored. A qualitative difference between the evolutionary dynamics of decision-making rules was observed in the work presented in Chapter 5 between finite and infinite groups, which can be interpreted to reflect this important distinction. By modeling the same pairwise strategy interactions for the two cases (i.e. finite and infinite groups) we can derive more robust results as well as explore the existence of a limit group size where a transition between the two kinds of dynamics are observed.

### 9.2.4 Bias Propagation

The preliminary analysis shown in Chapter 6 indicates that different decision-making rules lead to different long-term traits in the infinite group limit. It specifically demonstrates that the aggregate decision-making rule that leads do bias accumulation over time allows for the formation of information cascades (an event that occurs with a positive probability), while the simplified rule that does not accumulate bias in the same straightforward way will not lead to cascades.

The analysis here was focused on exploring the progression of decisions for the clear-cut case where the agents were observing information in regions of either $x > 0$ or $x < 0$. This has provided insight into the behaviour of a group under a stable social bias (i.e. social bias indicating only option $A$ or option $B$), while the deviation from this that was explored was that of the first return to the origin $x = 0$ after having spent time in $x \neq 0$. What is not addressed is the behaviour of a group that is transitioning between the regions of $x > 0$ and $x < 0$.

Overall, additional work is required for this to be robustly shown analytically since at this moment this result is weak. One specific future direction at this stage, is exploring the contradiction that arises from the long-term behaviour of the dynamic rule; while in Chapter 5 it is proposed that the value of $s$ of a group needs to be constrained for a group of finite size, in Chapter 6 it follows that in order to reach the optimal outcome ($A$ being chosen more frequently), the value of $s \to \infty$ as the group size gets larger. It is interesting to further examine this point, to better determine the relationship between group size and value of $s$.

Lastly, while it was possible to observe a connection between bias propagation and consensus decisions for the simplified versions of the model (namely the aggregate, the majority and the dynamic strategies), this was not done for the dependencies strategy. An equally simple model cannot be used for this strategy, given its more complex nature that does not permit the use of Markov chains as the decision-making is not memoryless; as a consequence finding long-term tendencies was not attempted outside of Chapter 4 where the strategy was explored using computational tools. It would however be interesting to explore a more analytical framework to describe the long-term tendencies of the dependencies strategy, and observe whether this is different or similar to any of the simplified version at the infinite size limit. Such an analysis may also show whether the dependencies strategy collapses to one of the simplified cases considered here, and if so under what conditions.

### 9.2.5 Optimisation

The work presented in Chapter 8 shows how sociality evolves in relation to an agent's place in the sequence (along with $a$ and strategy), and how this is always different to the collectively optimal one. A baseline case was explored here: the one where all other agents are employing the collectively optimal value of sociality $s_{opt}$ which is known from Chapter 5 to be unstable. This means that although the results of Chapter 8 are still informative and provide novel insight into those dynamics, some aspects of them were – to some extent – expected. Future work can explore the results of this study but for the case where all other agents in the group employ the evolutionary stable value of sociality $s_{ES}$ instead.

Additionally, as seen in Chapter 8, depending on $i$ each agent has a different distance from $s_{opt}$. This raises the question of whether there is a more complex strategy, that would still allow agents to use $s_i \neq s_{col}$ while at the same time keeping the optimal $\mathbb{E}(n_A)$ – or at least have it be close to its 'ideal' value. This question can be explored both for the cases where the group employs $s_{opt}$ and $s_{ESS}$.

# References

Andreoni, J., & Miller, J. H. (1993). Rational cooperation in the finitely repeated prisoner's dilemma: Experimental evidence. *The economic journal*, *103*(418), 570–585.

Aoki, I. (1982). A simulation study on the schooling mechanism in fish. *NIPPON SUISAN GAKKAISHI*, *48*(8), 1081–1088.

Aplin, L. M., Farine, D. R., Mann, R. P., & Sheldon, B. C. (2014). Individual-level personality influences social foraging and collective behaviour in wild birds. *Proceedings of the Royal Society B: Biological Sciences*, *281*(1789), 20141016.

Arganda, S., Pérez-Escudero, A., & De Polavieja, G. G. (2012). A common rule for decision making in animal collectives across species. *Proceedings of the National Academy of Sciences of the United States of America*, *109*(50), 20508–20513. https://doi.org/10.1073/pnas.1210664109

Axelrod, R., & Hamilton, W. D. (1981). The evolution of cooperation. *science*, *211*(4489), 1390–1396.

Bateson, M., Healy, S. D., & Hurly, T. A. (2002). Irrational choices in hummingbird foraging behaviour. *Animal Behaviour*, *63*(3), 587–596.

Beckers, R., Deneubourg, J. L., & Goss, S. (1993). Modulation of Trail Laying in the Ant Lasius niger (Hymenoptera: Formicidae) and Its Role in the Collective Selection of a Food Source. *Journal of lnsect Behavior*, *6*(6).

Bertsekas, D., & Tsitsiklis, J. N. (2008). *Introduction to probability* (Vol. 1). Athena Scientific.

Bialek, W., Cavagna, A., Giardina, I., Mora, T., Silvestri, E., Viale, M., & Walczak, A. M. (2012). Statistical mechanics for natural flocks of birds. *Proceedings of the National Academy of Sciences of the United States of America*, *109*(13), 4786–4791. https://doi.org/10.1073/PNAS.1118633109/SUPPL_FILE/APPENDIX.PDF

Biesmeijer, J. C., & Seeley, T. D. (2005). The use of waggle dance information by honey bees throughout their foraging careers. *Behavioral Ecology and Sociobiology*, *59*, 133–142.

Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of political Economy*, *100*(5), 992–1026.

Bode, N. W., Franks, D. W., Wood, A. J., Piercy, J. J., Croft, D. P., & Codling, E. A. (2012). Distinguishing social from nonsocial navigation in moving animal groups. *The American Naturalist*, *179*(5), 621–632.

Bode, N. W., Wood, A. J., & Franks, D. W. (2011). The impact of social networks on animal collective motion. *Animal Behaviour*, *82*(1), 29–38. https://doi.org/10.1016/j.anbehav.2011.04.011

Boland, P. J. (1989). Majority systems and the condorcet jury theorem. *Journal of the Royal Statistical Society Series D: The Statistician*, *38*(3), 181–189.

Brännström, Å., Johansson, J., & Von Festenberg, N. (2013). The hitchhiker's guide to adaptive dynamics. *Games*, *4*(3), 304–328.

Castellano, C., Fortunato, S., & Loreto, V. (2009). Statistical physics of social dynamics. *Reviews of modern physics*, *81*(2), 591.

Conradt, L. (1998). Could asynchrony in activity between the sexes cause intersexual social segregation in ruminants? *Proceedings of the Royal Society of London. Series B: Biological Sciences*, *265*(1403), 1359–1368.

Conradt, L., & Roper, T. J. (2000). Activity synchrony and social cohesion: a ssion-fusion model. https://doi.org/10.1098/rspb.2000.1271

Conradt, L., & Roper, T. J. (2003). Group decision-making in animals. *421*(January), 1996–1999.

Conradt, L., & Roper, T. (2007). Democracy in animals: the evolution of shared group decisions. *Proceedings of the Royal Society B: Biological Sciences*, *274*(1623), 2317–2326. https://doi.org/10.1098/rspb.2007.0186

Conradt, L., & Roper, T. J. (2009). Conflicts of interest and the evolution of decision sharing. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *364*(1518), 807–819. https://doi.org/10.1098/rstb.2008.0257

Cooper, R., DeJong, D. V., Forsythe, R., & Ross, T. W. (1996). Cooperation without reputation: Experimental evidence from prisoner's dilemma games. *Games and Economic Behavior*, *12*(2), 187–218.

Couzin, I. D., Krause, J., Franks, N. R., & Levin, S. A. (2005). Effective leadership and decision-making in animal groups on the move. *Nature*, *433*(7025), 513–516. https://doi.org/10.1038/nature03236

Couzin, I. D., Krause, J., James, R., Ruxton, G. D., & Franks, N. R. (2002). Collective memory and spatial sorting in animal groups. *Journal of theoretical biology*, *218*(1), 1–11.

Dall, S. R., Giraldeau, L.-A., Olsson, O., McNamara, J. M., & Stephens, D. W. (2005). Information and its use by animals in evolutionary ecology. *Trends in ecology & evolution*, *20*(4), 187–193.

Darwin, C. (1964). *On the origin of species: A facsimile of the first edition*. Harvard University Press.

Dawkins, R. (2016). *The selfish gene*. Oxford university press.

Deneubourg, J. L., Aron, S., Goss, S., Pasteels, J. M., & Duerinck, G. (1986). Random behaviour, amplification processes and number of participants: How they contribute to the foraging properties of ants. *Physica D: Nonlinear Phenomena*, *22*(1-3), 176–186. https://doi.org/10.1016/0167-2789(86)90239-3

Deneubourg, J.-L., Goss, S., Franks, N., & Pasteels, J. (1989). The blind leading the blind: Modeling chemically mediated army ant raid patterns. *Journal of insect behavior*, *2*, 719–725.

Dobzhansky, T. (2013). Nothing in biology makes sense except in the light of evolution. *The american biology teacher*, *75*(2), 87–91.

Dussutour, A., Beekman, M., Nicolis, S. C., & Meyer, B. (2009). Noise improves collective decision-making by ants in dynamic environments. *Proceedings of the Royal Society B: Biological Sciences*, *276*(1677), 4353–4361.

Eguíluz, V. M., Masuda, N., & Fernández-Gracia, J. (2015). Bayesian decision making in human collectives with binary choices. *PLoS One*, *10*(4), e0121332.

Faria, J. J., Codling, E. A., Dyer, J. R., Trillmich, F., & Krause, J. (2009). Navigation in human crowds; testing the many-wrongs principle. *Animal Behaviour*, *78*(3), 587–591.

Farine, D. R., Aplin, L. M., Garroway, C. J., Mann, R. P., & Sheldon, B. C. (2014). Collective decision making and social interaction rules in mixed-species flocks of songbirds. *Animal Behaviour*, *95*, 173–182.

Fawcett, T. W., Fallenstein, B., Higginson, A. D., Houston, A. I., Mallpress, D. E., Trimmer, P. C., & McNamara, J. M. (2014, March). The evolution of decision rules in complex environments. https://doi.org/10.1016/j.tics.2013.12.012

Fawcett, T. W., Hamblin, S., & Giraldeau, L.-A. (2013). Exposing the behavioral gambit: The evolution of learning and decision rules. *Behavioral Ecology*, *24*(1), 2–11.

Franks, N. R., Gomez, N., Goss, S., & Deneubourg, J.-L. (1991). The blind leading the blind in army ant raid patterns: Testing a model of self-organization (hymenoptera: Formicidae). *Journal of Insect Behavior*, *4*, 583–607.

Galam, S. (1997). Rational group decision making: A random field ising model at t= 0. *Physica A: Statistical Mechanics and its Applications*, *238*(1-4), 66–80.

Gautrais, J., Ginelli, F., Fournier, R., Blanco, S., Soria, M., Chaté, H., & Theraulaz, G. (2012). Deciphering interactions in moving animal groups. *PLOS Computational Biology*, *8*(9), e1002678.

Geritz, S. A., Kisdi, E., Mesze
NA, G., & Metz, J. A. (1998). Evolutionarily singular strategies and the adaptive growth and branching of the evolutionary tree. *Evolutionary ecology*, *12*, 35–57.

Giraldeau, L.-A., Valone, T. J., & Templeton, J. J. (2002). Potential disadvantages of using socially acquired information. https://doi.org/10.1098/rstb.2002.1065

Gómez-Laplaza, L. M., & Gerlai, R. (2011). Can angelfish (pterophyllum scalare) count? discrimination between different shoal sizes follows weber's law. *Animal cognition*, *14*, 1–9.

Gómez-Nava, L., Lange, R. T., Klamser, P. P., Lukas, J., Arias-Rodriguez, L., Bierbach, D., Krause, J., Sprekeler, H., & Romanczuk, P. (2023). Fish shoals resemble a stochastic excitable system driven by environmental perturbations. *Nature Physics*, 1–7.

Goss, S., Aron, S., Deneubourg, J.-L., & Pasteels, J. M. (1989). Self-organized shortcuts in the argentine ant. *Naturwissenschaften*, *76*(12), 579–581.

Grassé, P.-P. (1959). La reconstruction du nid et les coordinations interindividuelles chez bellicositermes natalensis et cubitermes sp. la théorie de la stigmergie: Essai d'interprétation du comportement des termites constructeurs. *Insectes sociaux*, *6*, 41–80.

Guttal, V., & Couzin, I. D. (2010). Social interactions, information use, and the evolution of collective migration. *Proceedings of the National Academy of Sciences of the United States of America*, *107*(37), 16172–16177. https://doi.org/10.1073/pnas.1006874107

Hamilton, W. D. (1964). The genetical evolution of social behaviour. ii. *Journal of theoretical biology*, *7*(1), 17–52.

Herbert-Read, J. E. (2016). Understanding how animal groups achieve coordinated movement. *Journal of Experimental Biology*, *219*(19), 2971–2983. https://doi.org/10.1242/JEB.129411

Herbert-Read, J. E., Perna, A., Mann, R. P., Schaerf, T. M., Sumpter, D. J. T., & Ward, A. J. W. (2011). Inferring the rules of interaction of shoaling fish. *Proceedings of the National Academy of Sciences*, *108*(46), 18726–18731.

Hinz, R. C., & De Polavieja, G. G. (2017). Ontogeny of collective behavior reveals a simple attraction rule. *Proceedings of the National Academy of Sciences*, *114*(9), 2295–2300.

Hofbauer, J., & Sigmund, K. (2003). Evolutionary game dynamics. *Bulletin of the American mathematical society*, *40*(4), 479–519.

Houston, A. I., & McNamara, J. M. (1999). *Models of adaptive behaviour: An approach based on state*. Cambridge University Press.

Houston, A. I., McNamara, J. M., & Steer, M. D. (2007). Violations of transitivity under fitness maximization. *Biology Letters*, *3*(4), 365–367.

Huepe, C., Ferrante, E., Wenseleers, T., & Turgut, A. E. (2015). Scale-Free Correlations in Flocking Systems with Position-Based Interactions. *Journal of Statistical Physics*, *158*(3), 549–562. https://doi.org/10.1007/S10955-014-1114-8/FIGURES/4

Hutchinson, J. M., & Gigerenzer, G. (2005). Simple heuristics and rules of thumb: Where psychologists and behavioural biologists might meet. *Behavioural processes*, *69*(2), 97–124.

Jayles, B., rin Kim, H., Escobedo, R., Cezera, S., Blanchet, A., Kameda, T., Sire, C., & Theraulaz, G. (2017). How social information can improve estimation accuracy in human groups. *Proceedings of the National Academy of Sciences of the United States of America*, *114*, 12620–12625. https://doi.org/10.1073/PNAS.1703695114/SUPPL_FILE/PNAS.1703695114.SAPP.PDF

Kadak, K., & Miller, N. (2020). Follow the straggler: zebrafish use a simple heuristic for collective decision-making. https://doi.org/10.1098/rspb.2020.2690

Katz, Y., Tunstrom, K., Ioannou, C., Huepe, C., & Couzin, I. (2011). Inferring the structure and dynamics of interactions in schooling fish. *Proceedings of the National Academy of Sciences*, *108*(46), 18720–18725.

King, A. J., & Cowlishaw, G. (2007). When to use social information: the advantage of large group size in individual decision making. *Biol. Lett*, *3*, 137–139. https://doi.org/10.1098/rsbl.2007.0017

Koutsoupias, E., & Papadimitriou, C. (2009). Worst-case equilibria. *Computer science review*, *3*(2), 65–69.

Krause, J., Ruxton, G. D., Ruxton, G., Ruxton, I. G., et al. (2002). *Living in groups*. Oxford University Press.

Kreps, D. M., Milgrom, P., Roberts, J., & Wilson, R. (1982). Rational cooperation in the finitely repeated prisoners' dilemma. *Journal of Economic theory*, *27*(2), 245–252.

Latty, T., & Beekman, M. (2011). Irrational decision-making in an amoeboid organism: Transitivity and context-dependent preferences. *Proceedings of the Royal Society B: Biological Sciences*, *278*(1703), 307–312.

Lecheval, V., Jiang, L., Tichit, P., Sire, C., Hemelrijk, C. K., & Theraulaz, G. (2018). Social conformity and propagation of information in collective u-turns of fish schools. *Proceedings of the Royal Society B: Biological Sciences*, *285*(1877). https://doi.org/10.1098/RSPB.2018.0251

Ling, H., Mclvor, G. E., van der Vaart, K., Vaughan, R. T., Thornton, A., & Ouellette, N. T. (2019). Local interactions and their group-level consequences in flocking jackdaws. https://doi.org/10.1098/rspb.2019.0865

Magurran, A., Oulton, W., & Pitcher, T. (1985). Vigilant behaviour and shoal size in minnows. *Zeitschrift für Tierpsychologie*, *67*(1-4), 167–178.

Mann, R. P. (2018). Collective decision making by rational individuals. *Proceedings of the National Academy of Sciences of the United States of America*, *115*(44), E10387–E10396. https://doi.org/10.1073/pnas.1811964115

Mann, R. P. (2020). Collective decision-making by rational agents with differing preferences. *Proceedings of the National Academy of Sciences*, *117*(19), 10388–10396. https://doi.org/10.1073/pnas.2000840117

Mann, R. P. (2021). Optimal use of simplified social information in sequential decision-making. *Journal of the Royal Society Interface*, *18*(179), 20210082.

Mann, R. P., Faria, J., Sumpter, D. J. T., & Krause, J. (2013). The dynamics of audience applause. https://doi.org/10.1098/rsif.2013.0466

Mann, R. P., Herbert-Read, J. E., Ma, Q., Jordan, L. A., Sumpter, D. J., & Ward, A. J. (2014). A model comparison reveals dynamic social information drives the movements of humbug damselfish (dascyllus aruanus). *Journal of the Royal Society Interface*, *11*(90), 20130794.

McNamara, J. M., & Houston, A. I. (2009). Integrating function and mechanism. *Trends in ecology & evolution*, *24*(12), 670–675.

Miller, N., Garnier, S., Hartnett, A. T., & Couzin, I. D. (2013). Both information and social cohesion determine collective decisions in animal groups. *Proc Nat Acad Sci USA*, *110*, 5263–5268.

Mora, T., & Bialek, W. (2011). Are biological systems poised at criticality? *Journal of Statistical Physics*, *144*, 268–302.

Munoz, M. A. (2018). Colloquium: Criticality and dynamical scaling in living systems. *Reviews of Modern Physics*, *90*(3), 031001.

Norris, J. R. (1998). *Markov chains*. Cambridge university press.

Nowak, M. (2019). *Evolutionary Dynamics: Exploring the Equations of Life*. https://doi.org/10.1017/9781108697385.002

Osborne, M. J., et al. (2004). *An introduction to game theory* (Vol. 3). Springer.

Parker, G. (1978). Selfish genes, evolutionary games, and the adaptiveness of behaviour. *Nature*, *274*(5674), 849–855.

Pérez-Escudero, A., & de Polavieja, G. G. (2011). Collective animal behavior from bayesian estimation and probability matching. *PLoS Computational Biology*, *7*(11), 14. https://doi.org/10.1371/journal.pcbi.1002282

Pérez-Escudero, A., & de Polavieja, G. G. (2017). Adversity magnifies the importance of social information in decision-making. *Journal of The Royal Society Interface*, *14*(136), 20170748.

Pérez-Escudero, A., Miller, N., Hartnett, A. T., Garnier, S., Couzin, I. D., & De Polavieja, G. G. (2013). Estimation models describe well collective decisions among three options. *Proceedings of the National Academy of Sciences*, *110*(37), E3466–E3467.

Perna, A., Granovskiy, B., Garnier, S., Nicolis, S. C., Labédan, M., Theraulaz, G., Fourcassié, V., & Sumpter, D. J. (2012). Individual rules for trail pattern formation in Argentine ants (linepithema humile). *PLoS Computational Biology*, *8*(7). https://doi.org/10.1371/journal.pcbi.1002592

Pettit, B., Perna, A., Biro, D., & Sumpter, D. J. T. (2013). Interaction rules underlying group decisions in homing pigeons. *Journal of the Royal Society Interface*, *10*(89), 20130529.

Pinkoviezky, I., Couzin, I. D., & Gov, N. S. (2018). Collective conflict resolution in groups on the move. *Physical Review E*, *97*(3), 032304.

Pothos, E. M., Perry, G., Corr, P. J., Matthew, M. R., & Busemeyer, J. R. (2011). Understanding cooperation in the prisoner's dilemma game. *Personality and Individual Differences*, *51*(3), 210–215.

Pratt, S. C., Mallon, E. B., Sumpter, D. J., & Franks, N. R. (2002). Quorum sensing, recruitment, and collective decision-making during colony emigration by the ant Leptothorax albipennis. *Behavioral Ecology and Sociobiology*, *52*(2), 117–127. https://doi.org/10.1007/s00265-002-0487-x

Preston, M. D., Pitchford, J. W., & Wood, A. J. (2010). Evolutionary optimality in stochastic search problems. *Journal of the Royal Society Interface*, *7*(50), 1301–1310.

Pulliam, H. R. (1973). On the advantages of flocking. *J. Theor. Biol.*, *38*(2), 419–422. https://doi.org/10.1016/0022-5193(73)90184-7

Reynolds, C. W. (1987). Flocks, herds and schools: A distributed behavioral model. *Proceedings of the 14th annual conference on Computer graphics and interactive techniques*, 25–34.

Rieucau, G., & Giraldeau, L. A. (2009). Persuasive companions can be wrong: The use of misleading social information in nutmeg mannikins. *Behavioral Ecology*, *20*, 1217–1222. https://doi.org/10.1093/BEHECO/ARP121

Romanczuk, P., & Daniels, B. C. (2023). Phase transitions and criticality in the collective behavior of animals—self-organization and biological function. In *Order, disorder and criticality: Advanced problems of phase transition theory* (pp. 179–208). World Scientific.

Romanczuk, P., & Schimansky-Geier, L. (2012). Swarming and pattern formation due to selective attraction and repulsion. *Interface Focus*, *2*(6), 746–756. https://doi.org/10.1098/RSFS.2012.0030

Selous, E. (1931). *Thought transference (or what?) in birds*. Constable; Company Ltd.

Shafir, S., Waite, T. A., & Smith, B. H. (2002). Context-dependent violations of rational choice in honeybees (apis mellifera) and gray jays (perisoreus canadensis). *Behavioral Ecology and Sociobiology*, *51*, 180–187.

Sigalou, A., & Mann, R. P. (2023). Evolutionary stability of social interaction rules in collective decision-making. *Physical Biology*, *20*(4), 045003.

Smith, J. M. (1984). Game theory and the evolution of behaviour. *Behavioral and Brain Sciences*, *7*(1), 95–101. https://doi.org/10.1017/S0140525X00026327

Smith, J. M., & Price, G. R. (1973). The logic of animal conflict. *Nature*, *246*(5427), 15–18. https://doi.org/10.1038/246015a0

Smith, J. M., & Parker, G. A. (1976). The logic of asymmetric contests. *Animal behaviour*, *24*(1), 159–175.

Sridhar, V. H., Li, L., Gorbonos, D., Nagy, M., Schell, B. R., Sorochkin, T., Gov, N. S., & Couzin, I. D. (2021). The geometry of decision-making in individuals and collectives. *Proceedings of the National Academy of Sciences*, *118*(50), e2102157118. https://doi.org/10.1073/pnas.2102157118

Strömbom, D. (2011). Collective motion from local attraction. *Journal of Theoretical Biology*, *283*(1), 145–151. https://doi.org/10.1016/J.JTBI.2011.05.019

Sumpter, D. J. (2006). The principles of collective animal behaviour. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *361*(1465), 5–22. https://doi.org/10.1098/rstb.2005.1733

Sumpter, D. J. T. (2010). *Collective Animal Behavior*.

Sumpter, D. J., Mann, R. P., & Perna, A. (2012). The modelling cycle for collective animal behaviour. *Interface focus*, *2*(6), 764–773.

Sumpter, D. J., & Pratt, S. C. (2009). Quorum responses and consensus decision making. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *364*(1518), 743–753. https://doi.org/10.1098/rstb.2008.0204

Taylor, P. D., & Jonker, L. B. (1978). Evolutionary stable strategies and game dynamics. *Mathematical biosciences*, *40*(1-2), 145–156.

Todd, P. M., & Gigerenzer, G. (2007). Environments that make us smart: Ecological rationality. *Current directions in psychological science*, *16*(3), 167–171.

Torney, C. J., Lorenzi, T., Couzin, I. D., & Levin, S. A. (2015). Social information use and the evolution of unresponsiveness in collective systems. *Journal of the Royal Society Interface*, *12*(103), 20140893.

Trimmer, P. C. (2013). Optimal behaviour can violate the principle of regularity. *Proceedings of the Royal Society B: Biological Sciences*, *280*(1763), 20130858.

Trivers, R. L. (1971). The evolution of reciprocal altruism. *The Quarterly review of biology*, *46*(1), 35–57.

Valone, T. J., & Templeton, J. J. (2002). Public information for the assessment of quality: a widespread social phenomenon. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, *357*(1427), 1549–1557. https://doi.org/10.1098/RSTB.2002.1064

Vicsek, T., Czirók, A., Ben-Jacob, E., Cohen, I., & Shochet, O. (1995). Novel type of phase transition in a system of self-driven particles. *Phys. Rev. Lett.*, *75*, 1226–1229. https://doi.org/10.1103/PhysRevLett.75.1226

Vicsek, T., & Zafeiris, A. (2012). Collective motion. *Physics Reports*, *517*(3-4), 71–140. https://doi.org/10.1016/J.PHYSREP.2012.03.004

Von Neumann, J., & Morgenstern, O. (2007). Theory of games and economic behavior: 60th anniversary commemorative edition. In *Theory of games and economic behavior*. Princeton university press.

Waite, T. A. (2001). Intransitive preferences in hoarding gray jays (perisoreus canadensis). *Behavioral Ecology and Sociobiology*, *50*, 116–121.

Ward, A. J., Sumpter, D. J., Couzin, I. D., Hart, P. J., & Krause, J. (2008). Quorum decision-making facilitates information transfer in fish shoals. *Proceedings of the National Academy of Sciences of the United States of America*, *105*, 6948–6953. https://doi.org/10.1073/pnas.0710344105

West, S. A., Griffin, A. S., & Gardner, A. (2007). Social semantics: Altruism, cooperation, mutualism, strong reciprocity and group selection. *Journal of evolutionary biology*, *20*(2), 415–432.

Wolf, M., Kurvers, R. H., Ward, A. J., Krause, S., & Krause, J. (2013). Accurate decisions in an uncertain world: Collective cognition increases true positives while decreasing false positives. *Proceedings of the Royal Society B: Biological Sciences*, *280*(1756), 20122777.

Wood, A. J., & Ackland, G. J. (2007). Evolving the selfish herd: Emergence of distinct aggregating strategies in an individual-based model. *Proceedings of the Royal Society B: Biological Sciences*, *274*(1618), 1637–1642.