# Decision making in a social world:
## Coordination, recommendation, and resource allocation

Dissertation submitted for the degree of
Doctor of Natural Sciences (Dr. rer. nat.)

presented by

**Nico Gradwohl**

at

Universität
Konstanz

Faculty of Sciences
Department of Psychology

Konstanz, 2021

Date of the oral examination: August 18, 2021
First referee: Prof. Dr. Wolfgang Gaissmaier
Second referee: Prof. Dr. Janina Hoffmann
Third referee: Prof. Dr. Thomas Müller

# Acknowledgments

This dissertation is the result of three and a half years of intensive work, that could not have been successful without the help of many important people. First I want to thank my main supervisor Wolfgang Gaissmaier. Thank you for many great years of supervision and collaboration in your working group, stimulating discussions, and your clear and direct (but never demeaning) feedback. I also want to thank my other two supervisors. I thank Janina Hoffmann for waking my interest in cognitive modelling, guiding me on the first steps, and teaching me to be critical and sceptical about virtually any result. Finally, I thank Thomas Müller for discussions of my work beyond the boundaries of disciplines with both him and his working group.

Next to my supervisors, I would like to thank Hansjörg Neth for extensive support, inspiring discussions, supporting comments, and his in-depth perspective on our joint projects. I also want to thank Helge Giese for his excellent guidance on multi-level modelling, as well as his always helpful support.

I also thank the other PhD students in our working group for great times at conferences, PhD-social meetings, and during organizing the JDMx workshop in Konstanz. Thank you, Kevin, Tamara, Nathalie, Tjaša, Sophie, and Ruchira! Special thanks to Felix, for pioneering a large part of conducting online group studies with oTree and being a great partner to share problems and solutions to the, sometimes exhausting, challenges of working with various software.

I am also grateful to the *Graduate School of Decision Sciences* (GSDS) for covering the largest period of my thesis with a scholarship and connecting me to a network of interdisciplinary PhD students. I would like to mention in particular Sandra, Theresa, and Ben for inspiring coffee breaks and broadening my horizon beyond psychological research to the political sciences. Moreover, I thank Sascha, Kevin, Michelle, and Mario for great times in the office. Although not precisely part of the GSDS, I want to thank Jana for being a great neighbor and friend.

Finally, I want to thank those outside the ivory tower of academia. Thanks to my parents for their support and interest in what I'm doing whatever I'm doing. Thanks also goes to my friends, including, among others, Uli, Max, Vincent, Lukas, Barry, and Mark for accepting that I essentially disappeared from their lives for some time. Especially I want to thank Fabienne. Without your patient support, an always open ear, nightly discussions, and endurance of the highs and lows I could never have finished this thesis. Most likely, there are many more people that I met along the way and that I could thank for contributing to this thesis in one or the other way. That said, I take for granted, that I have forgotten some people who I hope won't take any offense.

# Zusammenfassung

In dieser Dissertation wird beleuchtet, wie sich Menschen in einer sozialen Umwelt zurechtfinden, in der durch die Gegenwart anderer unterschiedliche, gegenseitige Abhängigkeiten entstehen. Ich untersuche wie Individuen soziale Information integrieren, wenn sie (1) sich unter Wettbewerb koordinieren müssen, (2) Ratschläge geben oder erhalten und (3) begrenzte Ressourcen zuteilen. Diese Situationen umfassen unterschiedliche Formen der Abhängigkeit, da in der ersten Situation die Belohnung des Einzelnen von der Gruppenentscheidung abhängt, in der zweiten Situation die Entscheidung durch die Information anderer beeinflusst wird und in der letzten Situation ein Individuum die Entscheidung über jemand anderen allein trifft. Fünf online-Experimente zeigen, dass Individuen in allen drei Situationen soziale Information mit eigener Information integrieren, um die Entscheidungssituationen zu meistern.

Der erste Forschungsaufsatz trägt zu einem besseren Verständnis darüber bei, wie Individuen lernen sich in einer Wettbewerbssituation zu koordinieren. Dabei kann es herausfordernd sein zu lernen, dass Ressourcen begrenzt sind und diese mit anderen geteilt werden müssen. Vermeintlich bessere Optionen können durch den Wettbewerb weniger attraktiv werden. Zwei Experimente zeigen, dass Individuen es schaffen ihr Verhalten zumindest teilweise an die Situation anzupassen. Überraschenderweise profitieren Individuen nur wenig davon, andere und deren Verhalten beobachten zu können. Es gibt aber dennoch Anzeichen dafür, dass bereits das bloße Wissen über die Anwesenheit anderer sowie die Möglichkeit andere zu beobachten mehr koordiniertes Verhalten ermöglichen. Insgesamt nutzen Individuen soziale Information zwar teilweise, verlassen sich aber stark auf ihre individuelle Information und ergänzen diese mit dem Wissen um die Anwesenheit anderer und deren Verhalten.

In großen Entscheidungsumwelten mit zahlreichen Optionen kann soziale Information besonders hilfreich sein. Jedoch stehen Individuen dabei immer vor dem “exploration-exploitation” Dilemma, ob sie bei einer guten Option bleiben wollen (ausbeuten) oder ob sie die Umwelt weiter erkunden möchten (explorieren), um eventuell eine bessere Option zu finden. Diese Entscheidung kann durch eine Empfehlung von außen beeinflusst werden. In zwei Experimenten haben wir untersucht, wie sich Geben und Erhalten einer Empfehlung auf Informationssuchverhalten in räumlich-korrelierten mehrarmigen Banditen auswirken. Das erste Experiment zeigt, dass das Geben einer Empfehlung die Informationssuche nur dann beeinflusst, wenn die Belohnung vom Erfolg des Empfehlungsempfängers abhängt. Die Umwelt wird dann weitläufiger exploriert. Im zweiten Experiment untersuchen wir, wie sich die Empfehlung auf die weitere Informationssuche des Empfängers auswirkt. Individuen werden von Empfehlungen angezogen. Es gelingt ihen dabei bedingt, Empfehlungen von vornherein auszuschließen, die sie aufgrund ihrer Vorerfahrung als minderwertig einstufen. Daneben explorieren sie bei hochwertigen Empfehlungen eher in deren näheren Umgebung weiter, während sie zu minderwertigen Empfehlungen größeren Abstand halten. Insgesamt finden wir, ähnlich wie im ersten Aufsatz, dass soziale Information, hier Empfehlungen, nicht in extremem Maße berücksichtigt werden.

Soziale Situationen umfassen nicht nur jene, in denen Entscheidungen in Abhängigkeit von anderen getroffen werden müssen, sondern auch solche, in denen andere vollständig von einer Entscheidung abhängen. Im dritten Forschungsaufsatz untersuchen wir hypothetische, und dennoch realistische, Triage-Entscheidungen während einer Pandemie. Individuen mussten entscheiden, welche*r von zwei Patient*innen das einzig verbleibende, lebensrettende Beatmungsgerät bekommen sollte. Wir untersuchen welche verschiedenen Eigenschaften der Patient*innen Individuen heranziehen, um ihre Entscheidung zu treffen. Unsere Ergebnisse legen nahe, dass für die Zuteilungsentscheidungen weitestgehend Richtlinien und utilitaristische Normen berücksichtigt werden. Allerdings integrieren einige Individuen auch Eigenschaften der Patient*innen mit eigenen Eigenschaften oder

ihrer eigenen Wahrnehmung über die COVID-19-Krise. Dabei bevorzugen Individuen jene Patient*innen, die ihnen ähnlich sind und retten eher jene, die sich an kooperative Normen halten.

Zusammengenommen heben diese drei Studien hervor, dass Individuen über Situationen mit unterschiedlichen Formen gegenseitiger Abhängigkeit hinweg soziale Information integrieren. Dabei diskutiere ich mögliche Prozesse, die der Integration sozialer und nicht-sozialer Inforomation zugrunde liegen. Dadurch, dass diese Dissertation wichtige Faktoren und Verbindungen zwischen scheinbar unterschiedlichen Situationen herstellt, trägt sie zu unserem Verständnis bei, wie sich Menschen in einer sozialen Umwelt zurechtfinden.

# Summary

This dissertation adds to our understanding how humans navigate a social world characterized by dependence and interdependence with others. More precisely, I investigate how individuals integrate social information, when they have to (1) establish coordination under competition, (2) search information when giving and receiving recommendations, and (3) allocate scarce medical resources. All situations involve different types of dependence, since decision outcomes are determined jointly, are affected through information of others, and one individual completely determines the outcomes of others. The results of five online experiments illustrate that individuals integrate social information with their own information to deal with the decision task across all three settings.

The first research paper adds to our understanding of how individuals learn to coordinate under competition. A crucial challenge can be to learn that resources are limited and need to be shared with others. Supposedly better options can lose attractiveness under competition. Although learning in situations of interdependence is challenging, two experiments show that individuals partially adapt their behavior. Surprisingly, individuals benefit little from observing how others behave, neither with, nor without interdependence. Yet, there is indication that already the mere knowledge of others' presence and observing others' behavior can provide an advantage for achieving more coordination. In sum, individuals strongly rely on their individual information but complement this information by observation of others or knowledge about their presence.

Social information can be especially helpful in large environments with large numbers of alternatives. Yet, individuals are confronted with the exploration-exploitation dilemma whether to stay with a good option or explore for a better one. Others can provide valuable information about the options in such an environment through giving a recommendation. In two experiments we investigated how providing and receiving of recommendations affect information search in spatially correlated bandits. Experiment 1 shows that information search among advisors is only affected if their rewards depend on the recipient. In that case they spread out their choices more. In Experiment 2 we investigate how receiving a recommendation affects information search among recipients. Recipients' exploration is initially attracted by recommended options. They partially succeed in filtering out low-valued recommendations based on their previous information. Moreover, they tend to be attracted by high-valued recommendations but repelled by low-valued recommendations. Nevertheless, similar to research paper 1 we find that individuals do not extremely rely on recommendations.

Social situations not only include those in which decisions have to be made under dependence from others, but also others may fully depend on a decision maker. In research paper 3 we investigated hypothetical, yet unfortunately realistic, triage decisions during a viral pandemic. We asked individuals to decide who of two should receive the only remaining ventilator to save their life. We investigate how individuals use various patient features to make their allocation decisions. Our results indicate that individuals' allocation decisions largely reflect guidelines and utilitarian norms. Yet, some individuals clearly integrate patient features with their own features, as well as perceptions of the current COVID-19 crisis. Individuals exhibit a bias for those who are similar to them and tend to save those who adhere to cooperative norms.

Taken together the three studies highlight that individuals integrate social information across scenarios that differ in the type of interdependence. Thereby, I discuss possible processes underlying the integration of social and non-social information. In highlighting important factors and connections between seemingly disparate situations the present thesis contributes to our understanding how humans make decisions in a social world.

# Contents

# List of Figures

# List of Tables

# 1

# Synopsis

## 1.1 Introduction

When did you make your last decision without any influence of others? Hard to tell, right? Others' influence on decision making is ubiquitous. Often decision outcomes may depend on others, others may provide us with information, and even the mere presence of others may affect our behavior.

Limited resources introduce competition, so that our rewards depend on the behavior of others, but others' rewards also depend on our actions. Remember when toilet-paper was suddenly scarce and you stood before the last package? You buying the package will let the next person go empty-handed, yet it was the behavior of others who brought you in this situation in the first place. We also depend on information we receive from others. This dissertation would not exist without the work that it builds upon or without what my colleagues, mentors, and supervisors taught me. Eventually, individuals may fully depend on others' decisions. Parents decide how much care to allocate to their offspring and doctors may decide on the best course of action.

The goal of this dissertation is to improve our understanding of how humans make decisions in a social world. But what makes a situation social? And which challenges do individuals face in a social world?

First, social situations are characterized by a dependence from others or their dependence on us. In what has been called the *minimal social situation*[1] (Kelley et al., 1962; Sidowski, 1957; Sidowski et al., 1956) individuals depend on each other without knowing about their interdependence and presence.

Beyond this bare minimum, many social situations involve varying degrees of social

---

[1] Not to be confused with the minimal *group* situation (e.g., Tajfel et al., 1971; Yamagishi & Kiyonari, 2000), in which categorization based on arbitrary features already introduces favoritism for those with whom one shares these features.

information (Frith & Frith, 2012), including observation of others and their experiences (Kendal et al., 2018; Mesoudi, 2016), as well as information about the identity (e.g., Tajfel et al., 1971; Tajfel & Turner, 1979) and the past behavior or reputation (Milinski et al., 2002; Nowak & Sigmund, 2005; Tennie et al., 2010) of social interaction partners. Thereby, one could even argue that social learning or copying (Kendal et al., 2018; Laland, 2004), or social influence processes (Cialdini & Goldstein, 2004; Deutsch & Gerard, 1955) *create* some kind of dependence, because the behavior of one individual influences the behavior of another individual. Because of their ubiquity, I will consider social situations that are characterized by varying degrees of interdependence and social information.

## 1.2 Challenges for decision making

In order to understand the challenges specific to decision making in a social world, it is important to understand which challenges the non-social world already creates. Generally, decisions can be considered successful, if they achieve satisfying goals (Dunwoody, 2009). Yet goal achievement can be complicated, because we often cannot know in advance, whether a certain choice or course of option will pay off. Decisions involve *risk* or *uncertainty* (Knight, 1921). Under risk, we have information about the probabilities, whereas under uncertainty there is no means of predicting the outcome.

Often, the only way to get information about available options is trial-and-error learning, based on whether an option is rewarding or not. Relatedly, in situations where options and their values are not known, individuals face an exploration-exploitation dilemma (see Cohen et al., 2007; Gershman, 2019; Hills et al., 2015; Mehlhorn et al., 2015). Should one try to find a superior option or should one stick with an already familiar and decent option? This kind of trade-off has been extensively investigated in multi-armed bandits (e.g., Lai & Robbins, 1985; Robbins, 1952; Steyvers et al., 2009; Sutton & Barto, 1998), where individuals have to decide between a number of options that are like slot-machines (or one-armed bandits—hence the name), which vary in their expected payoff. Thus, uncertainty and competing goals about searching and benefiting from information create one important challenge.

Another challenge is that, for any given decision, many pieces of information about options or multiple options can be available. To reach a decision between alternatives, individuals need to *integrate* pieces of information by weighing them or deciding which information to ignore (Lee & Cummins, 2004; Payne et al., 1988). Yet, they may benefit from ignoring some information, both because individuals' cognitive capacities may be limited and because ignoring part of the information can be beneficial (Gigerenzer & Gaissmaier, 2011). Thus, decision makers are confronted with the general challenges to determine which information to use, how to integrate it, and how to allocate their time to information search and the use of different options.

## 1.3 Decisions in a social world

On top of the challenges of non-social situations, social situations create additional challenges but also provide benefits. Overall, social situations typically constitute “large worlds” (Savage, 1954), in which multiple goals need to be considered at once (see Hertwig & Hoffrage, 2012). When outcomes depend on others, it can become even more difficult to predict, what a good choice is. Moreover, information provided by others can be helpful but also is yet another source of information that needs to be integrated.

### 1.3.1 Dependence and interdependence

As Kurt Lewin noted, interdependence is one large—if not the largest (cf. Marrow, 1969)—challenge of decisions in a social world. Individuals can depend on others in various ways. Managers make decisions regarding the fate of their employees, or citizens may vote on a policy and determine the further course of a society together. Similarly, consensus decisions can be described as unshared if one individual decides for the others, or as shared if the information from multiple individuals is combined (Conradt & Roper, 2005). Thus, individuals may depend on one another asymmetrically or jointly contribute to decision outcomes.

Thereby, one can distinguish different types of dependence based on *interdependence theory* (Balliet et al., 2017; Kelley & Thibaut, 1978). Outcomes of an actor can be affected by their own actions (*actor control*), the actions of an interaction partner (*partner control*), and the combination of both of their actions (*joint control*). Thus, in the extremes, a decision maker’s actions can fully determine the outcome of their partner(s), their outcomes can be fully determined by their partners’ actions, or all individuals’ actions can jointly determine the outcomes. Individuals can see themselves faced with dependence in either direction, or interdependence.

These extreme types of dependence all come with their own challenges for decision making. Situations of pure actor control are not social in itself and simply face the challenges for decision making described above. In situations of partner control other individuals’ outcomes depend on an actor’s decision. They raise questions how the other individual should be treated and how one determines their course of action. One may either aim to treat everyone equally or treat some individuals preferentially, for instance because of their group membership (Tajfel et al., 1971). Finally, interdependence or joint control imply that individuals need to determine their actions jointly to achieve the best possible outcome. Yet, *social dilemmas* (Dawes, 1980) and collective action problems (Ostrom, 1998) illustrate that this is challenging because individuals often have their own agenda. Making decisions in all three situations can be affected by various types of social information.

### 1.3.2 Social information

Beyond dependence and interdependence, a second crucial feature of social situations is social information. Just like dependence, social information can come in various types (e.g., Caldwell & Millen, 2009; Frith & Frith, 2012; Moussaid et al., 2009). Within the social learning literature social information typically refers to *observational information* about other individuals' behavior or outcomes (Kendal et al., 2018; Rieucau & Giraldeau, 2011). Yet social information can also take the form of *communication or teaching* (Aplin, 2019; Caldwell & Millen, 2009; Dean et al., 2014; Mesoudi, 2016). Even *indirect traces* that others' behavior leaves in the environment can serve as information, just like paths that others footsteps leave on the lawn (*stigmergy*, Moussaid et al., 2009; Theraulaz and Bonabeau, 1999). Eventually, social information can also be *about features of other individuals* (e.g., Frith & Frith, 2012).

Social learning by observing others can provide valuable information about the options in one's environment (Laland, 2004). In multi-armed bandits it has been extensively investigated, how individuals can benefit from observing and copying others if there are many options about which they would need to learn by trial and error learning (e.g., Efferson et al., 2008; Rendell et al., 2010; Toyokawa et al., 2014; Toyokawa et al., 2019). Results from a virtual tournament, which tested multiple strategies against each other, showed that strategies which rely on copying others are extremely successful to navigate such multi-armed bandits (Rendell et al., 2010).

Building upon others' experience allows to transcend the problem of limited time per individual to create cumulative achievements, like tools and solutions to problems that are part of culture in humans (Mesoudi, 2016) and non-human animals alike (Aplin, 2019). Social learning allows humans (as well as non-human animals) to achieve *cumulative cultural evolution* (see Boyd & Richerson, 1985, 1995; Dean et al., 2014). This means, that individuals combine their experience by learning from observing others or teaching about options and solutions. One famous example from the animal realm is that birds learn to solve problems, such as opening milk bottles (Sherry & Galef, 1984) or solving small tasks to open feeders (Aplin et al., 2015). Among humans, examples are uncountable and include low-level achievements like making increasingly useful tools (e.g., Mesoudi & O'Brien, 2008a), as well as high-level achievements like journeying into space, which required to combine a large amount of individual achievements (see Dean et al., 2014). Cultural evolution has also been modelled in the laboratory (for reviews see Caldwell et al., 2020; Miton & Charbonneau, 2018) through building paper planes (Caldwell & Millen, 2009) or (virtual) tools and innovations (e.g., Derex & Boyd, 2016; Mesoudi, 2011; Mesoudi & O'Brien, 2008a)

Beyond immediate effects of observation and teaching, also information from others who are not present can affect our behavior. Descriptive norms about what others do can affect individuals' environmental conservation behavior (Goldstein et al., 2008). Also mere perceptions about what others do can affect behavior, whether perceptions are accurate, as in the case of flu vaccination (Bruine de Bruin et al., 2019), or often biased, as in the case of alcohol consumption (Giese et al., 2019). Similarly, the mere presence of others can affect behavior and performance (for a review see Guerin, 1986). Thus, the influences of social situations range from very overt influences (like observation) to very subtle ones (like the presence of imagined others).

In addition to learning about the non-social environment from others and adapting our behavior in their presence, information about others' behavior and group membership allows us to determine our behavior towards them (see Everett et al., 2015; Frith & Frith, 2012). In cooperative interactions (with joint control) a strategy that takes into account past behavior leads to overall more cooperation (Axelrod & Hamilton, 1981). Likewise, individuals can build a reputation that they reciprocate cooperation, which others can use to determine whether one should cooperate with them or not (Milinski et al., 2002). Moreover, we may be more friendly and cooperative to those who are similar to us (e.g., Balliet et al., 2014, for a meta-analysis). *Social identity theory* (Tajfel & Turner, 1979) deals with the finding that sharing a group membership with others also strongly affects our responses to them (e.g., Tajfel et al., 1971). Reasons for this in-group favoritism could include higher expectations that they return the favor and a higher valuation of benefiting in-group members (Everett et al., 2015). This information about others helps individuals to successfully navigate interactions with others.

The potential of social information for helping to navigate social and non-social parts of our environment is immense. Yet, it is an additional source of information that individuals have to integrate. Essentially, social information can create a "trichotomy of exploit own knowledge–exploit others' knowledge–explore" (Frith & Frith, 2012, p. 295). Thus, despite benefits through additional information about the social and non-social environment, social situations involve additional challenges (see also van den Bos et al., 2013).

Not only do dependence and interdependence with others create challenges. Also the additional information that others provide needs to be integrated. And in fact, reliance on social information can render a situation of actor control social, because the outcome is now influenced by which information the other individual provided. Thus, both dependence and social information create challenges for making decisions in a social world.

The goal of the present thesis is to improve our understanding of how individuals make decisions in a social world. Therefore, we add to our understanding how individuals deal with situations that are characterized by different degrees of interdependence and social information. Obviously, both interdependence and social information have received considerable attention across fields and disciplines. Yet, it is crucial to investigate how individuals navigate social settings from different angles. We ask, how individuals rely on different sources of information to navigate settings with different dependence structures. This way we can uncover commonalities and differences, which ultimately can inform models of human decision making in social world. In this thesis I investigate 3 challenges. Specifically I address, how humans decide if they (1) have to establish coordination under competition, (2) search information when giving and receiving recommendations, and (3) allocate scarce medical resources.

## 1.4 Coordination under competition

Coordination is one crucial task that individuals in social contexts face. Several examples show a remarkable success at coordinating behavior between large numbers of individuals, among humans and non-human animals alike (Conradt & Roper, 2005; Couzin, 2009;

Moussaid et al., 2009). Remarkable examples of successful coordination include self organization phenomena like lane formation in pedestrians (e.g., Helbing et al., 2005; Helbing et al., 2001) and ants (e.g., Couzin & Franks, 2003), reaching consensus about where to forage in honeybees (Seeley, 2003), or where to travel in migratory birds (Black, 1988), as well as individuals distributing almost optimally in foraging-type situations like food search in animals (see Kennedy & Gray, 1993), but also panhandling children (Disma et al., 2011). Yet, a huge number of failures to adjust behavior according to others highlights the challenging nature of establishing coordination. Think of tackling climate change, which requires individuals to contribute by changing behavior and habits to achieve a greater goal (see Newell et al., 2014). Such collective action problems (Ostrom, 1998) are extremely difficult because individuals have goals which conflict with a common objective. And even if all individuals have the same goal, obstacles like limited communication structures (e.g., Shirado & Christakis, 2017) and uncertainty about others' behavior (e.g., Huyck et al., 1990; Rapoport et al., 2002) can make coordination a challenging task.

One factor that creates this kind of challenges is that resources are limited. Limited resources create interdependence in the form of competition and the need for coordination. Foraging individuals should optimally distribute according to the quality of resources, that is, they should form an *Ideal Free Distribution* (IFD, Fretwell and Lucas, 1969; Kennedy and Gray, 1993). Their joint actions determine the outcome of each individual (i.e., the share that each individual receives). Only those groups which distribute according to resource quality, will provide each individual with the maximum possible share of the resources, whereas over- or under-exploitation of a resource leads to lower outcomes for some individuals (and potentially to higher outcomes for others, see also Gallistel, 1990).

Crucially, the behavior to achieve a goal best can vastly differ between situations with and without interdependence. In situations of independence all individuals should simply try and find the option that is most rewarding. However, when limited resources create (inter)dependence the same option may be subject to fierce competition. Think of choosing your way to work, where one road may be clearly faster and preferable, yet becomes the slower choice through traffic. Similar situations have been investigated in congestion games (e.g., Selten et al., 2007). A generally less rewarding option may yield a higher prospect of yielding any reward or promise that the (smaller) reward has to be shared with fewer others. Accordingly, an individual may select the smaller road that is less frequented. Thus, interdependence can make a supposedly inferior option superior overall.

A key challenge for successful coordination is that it may require multiple processes. Individuals may simultaneously need to learn about an uncertain environment and about being in a situation of interdependence. Whereas we know where we want to go and somehow learned (typically not first-hand) that bumping into a stranger is an unpleasant experience, in some situations we need to learn what a desirable option is and whether its desirability is affected by others. Often options do not come pre-labelled with what is good and what is bad, but their quality has to be learned from experience (Hertwig et al., 2004; Hertwig & Erev, 2009). Thus, whenever we enter a new environment (e.g., after moving to a new place) we may need to learn about the value of available options and whether we compete with others for them. In addition to that, risk and uncertainty can pose a challenge to learning. The outcomes from decisions may not always be the same but fluctuate. Thus, successful coordination may require to learn at once, which options are good and whether one competes with others for these options. Thereby, fluctuations in decision outcomes may both originate from the environment or the behavior of others.

Moreover, successful coordination requires some kind of social information. Although observing others and their behavior (e.g., other pedestrians) is obviously helpful for coordination, self-organization and coordination processes may even work with very limited social information, including others' traces in the environment (Moussaid et al., 2009). Likewise, others' traces could also reflect in an individual's rewards, like when repeatedly encountering emptied or diminished resources. In fact, humans in a group foraging task still approximated coordination when they could not observe others' behavior (Goldstone & Ashpole, 2004; Goldstone et al., 2005). This implies that the social information required for successful coordination may be as subtle as changes in rewards which hint at the presence of others.

Each of the above challenges for coordination has been investigated in previous research. Yet they have not been investigated in concert, disentangling effects of the underlying interdependence and social information. Results from experiments in Behavioral Game Theory show that people can coordinate in various situations of interdependence (Camerer, 2003). Research on learning (see e.g., Erev & Barron, 2005; Estes, 1976; Gershman, 2019; Mehlhorn et al., 2015; Sutton & Barto, 1998) and decisions from experience (Hertwig et al., 2004; Hertwig & Erev, 2009) addressed how individuals cope with environments in which option values or probabilities need to be learned from experience. Research on group foraging clearly shows that groups can distribute approximately optimally when resources are limited (Goldstone & Ashpole, 2004; Goldstone et al., 2005; Kraft & Baum, 2001; Kraft et al., 2002; Sokolowski et al., 1999). In fact, humans can even establish some degree of coordination when they are unaware of their interdependence (Kelley et al., 1962). What is less clear, however, is how groups of individuals can establish coordination, if both, the value of choices and the interdependence with others have to be learned.

### 1.4.1 Summary of research paper 1 on coordination under undisclosed competition

One instance of learning to make decisions under risk and uncertainty is probability learning (Estes, 1976). In a typical task, an individual has to select between two mutually exclusive and exhaustive options one of which is rewarded with a higher probability than the other. Although the optimal strategy is to exclusively select the more likely option (*probability maximizing*), individuals often select the options approximately relative to their probability of being rewarded (*probability matching*). Are people simply irrational or limited in their cognitive abilities (e.g., Koehler & James, 2009)? Whereas the tenets of classical rational choice theory prescribe that individuals should optimize their rewards by maximizing, the answer is relatively clearly no.

In fact, there may be reasonable explanations, why individuals fall short of probability maximizing. It has been argued that individuals' goal is to outmaneuver the environment (Arkes et al., 1986; Fantino & Esfandiari, 2002) and it could be shown that matching probabilities can actually help to detect patterns (Gaissmaier & Schooler, 2008; Wilke et al., 2014). Moreover, in competitive situations with limited resources, matching the probabilities can be more optimal than maximizing. This avoids that all individuals compete for the rewards of the more likely option (Gallistel, 1990). In fact, when the number of individuals is large, their behavior actually results in an optimally coordinated distribution (Thuijsman et al., 1995). Thus, contrary to tenets of classical rational choice theory, one should not necessarily aim to identify and exploit the option that is most likely rewarding when competing for limited resources.

In line with such an adaptive value of probability matching, individuals' adaptively adjust their behavior towards probability matching in competitive interactions (Schulze et al., 2015). Thus, strategies related to probability matching may be a robust way to tackle situations in which competition may be ambiguous. Similarly, dyads have been shown to distribute their choices to both options to optimally exploit a probability learning task, but only, when given an additional incentive to increase group-level rewards (Schulze & Newell, 2015). However, these studies do not answer how individuals learn to coordinate in a setting with undisclosed competition, because individuals always knew at least about each others presence.

We ask how small groups of individuals learn to establish coordination under undisclosed competition. Here I briefly summarize the research presented in Chapter 2. Groups of 3 players played 200 rounds of a probability learning task. The more likely option was rewarded twice as likely ($p = 2/3$) as the less likely option ($p = 1/3$). To establish successful coordination in the competitive setting, individuals must find a way to distribute themselves to the options according to an IFD, that is 2 individuals in the more likely and 1 individual in the less likely option.

We conducted two experiments. In the first experiment, 1 human player played together with 2 reinforcement learning agents and in the second experiment 3 human players played the task together. To investigate how individuals learn to coordinate under risk and uncertainty, we contrasted competitive and non-competitive situations. Either individuals' rewards were independent of others' actions or rewards were split among all who selected the rewarded option. In non-competitive situations the optimal behavior is to exclusively select the more likely option, whereas in the competitive situation individuals should optimally distribute according to the option probabilities—that is, 2 individuals selecting the more likely, and 1 individual selecting the less likely option. We additionally varied, whether individuals were not informed (unaware) or informed (aware) about the presence of others and whether they could observe others' players behavior and rewards (observation).

This design allows us to disentangle the effects of the competitive reward structure and social information. In fact, individuals may rely on relatively robust strategies that result in probability matching to achieve a decent degree of coordination. Yet, they also may simply maximize or aim to achieve stable coordination.

Across both experiments we find that individuals adapt their behavior to the competitive reward structure by selecting the option that is rewarded more likely with a lower probability than if they do not compete. The likelihood of forming a coordinated state in Experiment 2 is increased in the competitive setting. Yet, not a single group established stable coordination in the last 50 rounds. Different levels of social information did not differ from each other.

Nevertheless, we do find evidence that the available social information affects coordination in different ways. Individuals who could observe each other established a coordinated (2:1) state more likely in competitive than in non-competitive settings and did so more likely than a group of individuals each of which behaves randomly. The differences between competitive and non-competitive settings were smallest when individuals were aware of each other. This could indicate that the mere information about others' presence can make individuals more likely to interpret the situation as competitive and adapt their behavior. In line with this idea, the difference between competitive and non-competitive settings was larger among groups in which all individuals had higher levels of Cognitive Reflection (Frederick, 2005). This could indicate that some individuals rely on robust strategies, or that they fail to override an initial competitive representation.

These findings suggest that individuals adapt to undisclosed interdependence but do so incompletely. Although the competitive nature of the task appears to influence behavior most, both competition and social information shape group-level outcomes. In fact, observation of others facilitates coordination to a small degree. Apparently, individuals integrate individual reward information with social information to inform their decisions.

### 1.4.2 Implications and future directions

The findings make apparent that learning to coordinate under risk and uncertainty is challenging. But why don't individuals mange to establish stable coordination? A comparison to other settings which allow more successful coordination helps to elucidate this.

Accordingly, we can identify three main challenges. First, uncertainty about the options requires individuals to learn what their goal is. Second, individuals were unaware of their interdependence and, third, could not observe each other in some conditions. In contrast, pedestrians, for instance, typically know where their goal is, are well aware of their interdependence, and can observe each others, and thus, only face uncertainty about the other individuals' next move. Similarly, in other group foraging experiments (Critchfield & Atteberry, 2003; Goldstone & Ashpole, 2004; Goldstone et al., 2005; Madden et al., 2002; Sokolowski et al., 1999) individuals were informed about their interdependence. Yet, they still approached coordination when they could not observe each other and had to learn about the quality of resource pools (Goldstone & Ashpole, 2004; Goldstone et al., 2005), although this proved difficult when resources were distributed probabilistically (Kraft et al., 2002). This indicates that one main difficulty for coordination lies in uncovering the interdependence structure under risk and uncertainty.

Although coordination is challenging, both interdependence with others and social information jointly shape learning in decisions under risk and uncertainty. Generally speaking, individuals apparently integrate information about their own rewards with information about their social surroundings. But how do interdependence and social information shape decisions?

To some degree social information allows to tackle interdependence. A coordinated state is most likely achieved if individuals can observe each other. Therefore, an explanation why individuals hardly achieved more coordination when they could observe each other than when they did not, may be that they used social information too little. That individuals rely less on social information than optimal is generally in line with previous research on social learning (e.g., Acerbi et al., 2016; Efferson et al., 2008; Mesoudi, 2011). Crucially, individuals may be more likely to copy when social information helps them to navigate an uncertain environment (for reviews see Kendal et al., 2018; Laland, 2004). Thus, under competition strategic uncertainty about others' behavior may have led social information to be dismissed to be of limited value.

Beyond that, the present results open the intriguing possibility that social information changes the way in which the individual reward signal is interpreted. Similar to indirect information transmission processes (Heylighen, 2016; Theraulaz & Bonabeau, 1999), individuals may interpret changes in their rewards as traces of others' behavior. Individuals may consider non-rewarded rounds as a result of competition and, therefore, process rewards by using strategies adapted to competition, even if resources are not limited. This relates to other situations in which individuals are affected by the perceived behavior of others like descriptive social norms (see Cialdini & Trost, 1998; Cialdini & Goldstein, 2004) or the mere presence of others (see Guerin, 1986). Like a representation of what a majority of others does affects behavior, like alcohol consumption (Borsari & Carey, 2003; Giese et al., 2019), a representation of interdependence might change the interpretation of one's reward signal.

How does the above translate into hypotheses about the underlying individual-level processes? Responses to interdependence in the absence of social information suggest that behavior is likely well described by reinforcement learning processes (Gershman & Daw, 2017; Sutton & Barto, 1998), which have been implied in the context of probability matching (Feher da Silva et al., 2017). Accordingly, individuals use their own rewards to update the expected value of the options, based on which they update their choice probabilities. The influences of social information indicate that this model may be modified by the observation of others. In line with this idea, Schulze et al. (2015) found that a model that rewarded whether one's choices diverged from those of a competitor described individuals' responses to a competitor well. Yet, based on our results one may need to add that rewards may be processed differently when individuals know of others' presence.

However, it has been argued that this kind of model remains more of an "as-if" description, than describing the actual process (Gigerenzer, 2020). An alternative are (heuristic) strategies which describe the processing steps. Such strategies could begin from simple win-stay, lose-shift strategies (Kelley et al., 1962; Nowak & Sigmund, 1993), according to which individuals switch options, as soon as they are not rewarded. The spectrum ends at more sophisticated strategies, which incorporate information about real or imagined others. Information about observable others could be included by incorporating rewards for avoiding them (Schulze et al., 2015). Additionally, information about imagined or non-observable others could affect the way in which rewards are processed or decisions are made. Models could reflect this, by weighting unrewarded rounds or decreases in rewards so, that they lead to more random behavior. Future work should aim to get a more accurate description of these underlying processes by competitively testing process models and testing them experimentally.

## 1.5 Recommendation and exploration

Coordination can benefit from social information. Although, the benefits can be limited in situations which additionally require learning about interdependence. The finding that individuals benefit little from social information inspired a closer investigation of how individuals rely on social information under various circumstances. Specifically, beyond challenges of learning and interdependence, the previous section already involved a second challenge: individuals have to integrate information provided by others with their own information. Thereby, social information comes in various shapes and sizes, including not only the observation of others but also teaching or advice and recommendations (see Caldwell & Millen, 2009).

Although copying others allows us to transcend our own limited experience, it creates dependence from others and can change a situation of actor control to one of partner control or joint control. Essentially, the outcomes of the copying individual are fully determined by the partner they copy (*partner control*). Yet, this dependence is preceded by a decision, how much to rely on this social information. Given that this reliance is under the control of the copying individual they first exert *agent control*, creating a situation of interdependence overall. Thus, relying on social information creates dependence only if a decision maker decides to do so.

Problems of strong dependence from others through copying arise because past errors may be propagated and amplified (Rieucau & Giraldeau, 2011). If we were to always copy those who came before us, we might miss out on changes in the environment and novel innovations (Boyd & Richerson, 1995; Rogers, 1988). Outdated information may be propagated (Kameda & Nakanishi, 2002; Laland & Williams, 1998), as well as maladaptive trends (Bikhchandani et al., 1998; Giraldeau et al., 2002). The propensity to copy others can lead to herding behavior (Kameda & Hastie, 2015; Raafat et al., 2009). This can lead to phenomena that are hardly more than a nuisance, like in the case any song can be made popular, irrespective of its quality (Salganik et al., 2006), but also can be related to far-reaching outcomes, like financial crisis (e.g., Chari & Kehoe, 2004). Thus, relying on others can create “collective wisdom” or “collective madness” (Toyokawa et al., 2019), dependent on how individuals rely on the information. Put bluntly: if everyone in a group relied on the experience of others, everyone would end up doing the same, whether for good or for bad.

But how do individuals resolve the decision about how much they depend on social information? On the one hand, individuals seemingly overuse information provided by others and produce informational cascades leading to financial crashes and other catastrophes. Informational cascades and other real-world examples suggest that humans rely strongly on social information. Similarly, humans often fail to ignore advice, even if it comes from an inexperienced advisor (Harvey & Fischer, 1997) or is of low quality or even random (Fiedler et al., 2019; Hütter & Fiedler, 2019; Schultze et al., 2017).

On the other hand, individuals often rely less on social information than they should. When individuals integrate their own information with social information, they often rely more strongly on individual trial-and-error learning and forego the potential benefits of copying others (Acerbi et al., 2016; Efferson et al., 2008; McElreath et al., 2008; McElreath et al., 2005; Mesoudi, 2011; Mesoudi & O'Brien, 2008a). Similarly, individuals show a robust tendency for *egocentric discounting* (Bonaccio & Dalal, 2006; Rader et al., 2017; Yaniv & Kleinberger, 2000). That is, they adjust their own estimate of some quantity (e.g., the length of a river) insufficiently after they receive advice. Thus, when individuals receive advice, they give more weight to their own information. It seems puzzling, why we observe both maladaptive herding and a frequent underuse of the available social information relative to individual information.

One important factor seems to be uncertainty. When their own information is uncertain, individuals rely more strongly on social information (Laland, 2004). Besides options which are themselves unpredictable also a large number of options creates uncertainty about which option is good and which is not. Since individuals often cannot explore all available alternatives, they need to allocate their time to searching and harvesting the fruits of good alternatives they already discovered. Essentially, they face the exploration-exploitation trade-off (Cohen et al., 2007; Mehlhorn et al., 2015). That is, they need to decide, whether to search and try out new options or stick with a good option they already found.

Thereby, the search space is often multi-dimensional and multi-modal (e.g., Acerbi et al., 2016; Derex & Boyd, 2016; Mesoudi, 2011; Mesoudi & O'Brien, 2008a). That is, there is a number of dimensions that can be varied and multiple solutions can provide good results. Thus, these environments also feature local optima, from which a solution cannot be improved further by minor changes. Previous studies relied, for instance, on building paper planes (Caldwell & Millen, 2009) or virtual projectiles (Mesoudi & O'Brien, 2008a). Thus, the exploration-exploitation dilemma reflects in the decision whether to try out different solutions or to stick to a well-working solution, at the risk of getting stuck in a local optimum.

Social information, like observation of others' behavior, but also recommendations, can provide valuable guidance to identify good or even the best option and allows groups to explore a larger number of options. Information search can happen on two sides of the exchange of social information and both need to decide about exploration and exploitation: On the one hand, the individual that provides others with social information needs to get that information through trying out solutions themselves or copying them from others. On the other hand, the individual who receives social information needs to decide, how much to depend on this information. Thus, they face the "trichotomy of exploit own knowledge–exploit others' knowledge–explore" (Frith & Frith, 2012, p. 295).

How this dilemma is resolved affects whether individuals improve their performance over time. As soon as multiple individuals decide to exploit information obtained from their predecessor, improvements will come to a halt. We contribute to our understanding of these transmission dynamics by investigating how providing and receiving recommendations affect information search. I present 2 experiments on how giving and receiving recommendations affect information search in Chapter 3; that I summarize in the following.

### 1.5.1 Summary of research project 2 on giving and receiving recommendations

Advice or recommendations are one frequent form of social information provided by others. Another, typically more experienced, individual may provide information to another individual (Bonaccio & Dalal, 2006). At least in this case, the interdependence of a social information recipient and a provider of social information need not be unidirectional. A financial advisor may be held accountable for the outcome their advice generates (Dana & Cain, 2015) and face retribution, if their advice turns out to be bad. Thus, the advisor's outcome may be (at least partially) determined by the recipient of this information.

Recipients of recommendations, often have previous information at their disposal, that allows them to judge the information they receive. Therefore, we investigated advice giving and receiving in *spatially correlated bandits* (Schulz et al., 2019; Wu et al., 2018). This kind of task reflects situations with a large number of options and multiple modes in which individuals can generalize previous experience to unknown options. Options are correlated, so that good options are more likely to be close to good and bad options are more likely close to bad options. Therefore, individuals can generalize from previous experience to guide their exploration and focus on the surroundings of high-valued options.

First, we investigated the advisor as the first part of the link. Mainly, we addressed the question whether the mere prospect of providing a recommendation changes how advisors search information. The mere prospect of providing a recommendation did not alter their information search relative to a condition where they were rewarded for accumulating points for themselves. Only if advisors were rewarded according to recipient performance, their information search became similar to their information search, when they were rewarded for finding the best option. Thus, interdependence alters advisors' search behavior.

Second, we investigated how the recipients of recommendations integrate recommendations with their previous experience. In fact, it has been argued that advice under uncertainty may be intrinsically rewarding (Biele et al., 2009; Biele et al., 2011). A first set of simulations shows, that exploiting recommendations increases cumulative rewards, especially if the recommendation was provided by an experienced advisor. Our experiment on advisors' information search indicated that a key simulation assumption that advisors pass on their best option was mainly warranted. Since a bias towards recommendations could be adaptive, we asked how it would shape recipients' exploration behavior, when a recommendation was received.

Simulations showed, that a combination of generalization with an additional bias for recommendations can help to avoid recommendations about which nearby known options provide information. Yet, this bias could lead individuals to resolve the exploration-exploitation dilemma rather towards exploitation, because this leads to an increased propensity to click on the recommendation. This comes at the cost of a decreased chance of finding the best option in the environment as compared to continued exploration. Given that individuals rely on the same process, they could accept an increased risk of getting stuck in inferior options, to both avoid bad recommendations and to increase their accumulated rewards.

Based on these insights, we experimentally investigated how recipients change their exploration behavior when they receive a recommendation. To address the question whether individuals successfully generalize from their previous experience, we manipulated recommendations in a 2×2-within-subjects design. Recommendations were always options that participants had not selected previously. We varied whether they were familiar (i.e., in close proximity to already revealed options) or novel (i.e., in some distance from known options) and whether their value was rather high or low. Simulations had indicated that individuals can apply generalization to familiar, but not to novel recommendations. Like in the simulations their goal was to accumulate rewards.

Additionally, we systematically varied recipient experience, resulting in a $2 \times 2 \times 2$-design. We conjectured, that less experienced individuals might be more prone to get stuck in inferior recommendations. Finally, we randomly varied the predictability of the environment through the degree of correlation between tiles in the environment in two stages (smooth and rough). We conjectured that environment predictability could affect the use of recommendations, since it was shown to affect exploration in previous studies of spatially correlated bandits (Wu et al., 2018).

We find, that exploration behavior is strongly affected by receiving a recommendation. Individuals overall have a high likelihood of clicking on a recommendation. A preference for recommendation familiarity in smooth environments is reminiscent of egocentric discounting (Yaniv & Kleinberger, 2000). That this effect is largely absent in rough environments may indicate that individuals rely more on novel information if the environment is less predictable. This is in line with the idea that individuals rely more on social information if the environment is more uncertain (Laland, 2004). In line with the idea that individuals generalize, we find evidence that individuals in predictable environments more likely select high-valued familiar recommendations than low-valued familiar recommendations. Thus, individuals may rely on adaptive processes to dismiss recommendations but, nevertheless, rarely ignore them completely.

Recommendations have an impact on exploration beyond leading individuals to click on them. First, individuals explore closer to higher-valued recommendations. Moreover, an increased distance to low-valued recommendations indicates that individuals also negative information to inform their further exploration. Essentially, bad recommendations could create a repulsion effect, leading individuals to avoid the surroundings of bad recommendations.

Yet, evidence that recommendations are treated differently from self-selected options is mixed. Large parts of the effect on exploration behavior are not sustained, after the recommendation is received. We do not find evidence that individuals generally get stuck in options inferior to the global optimum. Generally, individuals rarely cease exploration, whether they receive a recommendation or not. In fact, they do not exploit recommendations, although they could have been better off if they did. However, the repulsion effect persists, when recommendations are compared to self-selected options of the same value.

Moreover, we find results suggesting that the conditions for getting stuck are met among less experienced individuals in more predictable environments. They are less likely to find the global optimum, if they receive a high-valued recommendation than if they find a high-valued option themselves. Relatedly, individuals are less likely to find the globally optimal option, when they receive a locally optimal recommendation. Thereby, we do not find evidence that locally optimal recommendations differ from self-selected local optima.

In sum, providing and receiving advice can both affect information search. Among advisors this effect was limited to situations in which their rewards depend on the recipient. Among recipients both an attraction to unknown recommendations and their subsequent integration affect exploration behavior. Although it is a viable strategy, recipients thereby typically do not cease exploration after they receive a recommendation. This avoids that they get stuck in options inferior to the global optimum, although further exploration is costly on average. Overall, the information search of both advisors and recipients is affected by social information.

### 1.5.2 Implications and future directions

Chapter 3 addresses how the recommendation situation shapes information search in multi-modal environments. When advisors' rewards depend on recipients' success they alter their exploration behavior to search less locally, essentially exploiting their previous experience less. Further, it provides an investigation of how individuals decide to rely on recommendations, when this creates dependence. Among recipients, receiving social information in form of a recommendation clearly affects their exploration behavior, leading them towards recommended options. Yet, although recommendations in large multi-modal search environments could lead individuals to get stuck in inferior options, individuals rarely accept the complete dependence that would let them get stuck. In line with previous research, individuals appear to rely less on social information than optimal (Efferson et al., 2008; Mesoudi, 2011; Morgan et al., 2012). Thus, although the social situation affects exploration behavior, its effects among advisors are limited to situations of interdependence and the impact on recipients does not appear to be sustained in most conditions.

In fact, recipients' comparably weak reliance on recommendations could be beneficial on the level of groups and collectives. As soon as individuals cease to acquire novel information, social learning can become inferior and collective progress towards better solutions may be stalled (Rogers, 1988). Similar processes may also result in maladaptive informational cascades (Bikhchandani et al., 1998). Yet, Chapter 3 suggests that individuals who have the capacity to acquire individual information will also do so, even despite potential benefits from relying on recommendations from experienced advisors. Consequently, humans' stable tendency for the egocentric discounting of advice and recommendations (Bonaccio & Dalal, 2006) may not only serve to avoid individually harmful outlier opinions (Harries et al., 2004). Beyond the individual level, egocentric discounting may be related to continued information search, which can help groups and collectives to avoid getting stuck.

An explanation, why individuals do not cease exploration is that individuals are motivated by resolving uncertainty about unknown options (Wu et al., 2018). In fact, this could also be an additional explanation why achieving coordination in uncertain environments is challenging in Chapter 2. Since individuals aim to resolve uncertainty, they fail to stably select an option. Hence, the motivation to resolve uncertainty may not only be individually costly but also may hinder coordination.

The finding that individuals appear motivated to resolve uncertainty and continue exploration also has important implications for transmission chains that extend beyond a single dyadic link. Such transmission chains have been investigated in previous studies (e.g., Acerbi et al., 2016; Mesoudi & O'Brien, 2008a; Mesoudi & Whiten, 2008; Yahosseini & Moussaïd, 2020). Their results indicate that the collective performance over time can benefit from the transmission of information. Our results corroborate these findings in the context of recommendations and suggest that especially novel recommendations may help to find the best option. This underscores the importance of novel (Rogers, 1988) and uncorrelated (Lorenz et al., 2011) information for collective performance.

Beyond that, our results suggest that individuals will likely be able to provide novel information to the collective if they have experience of their own and resources to explore the environment further. This implies, that it is not to be expected that transmission chains with more than one dyadic link get indefinitely stuck in inferior options. Yet, individuals with little experience may slow down the process in predictable environments. In fact, recipients' resources may play a crucial role for the absence of detrimental effects of recommendations. Similarly, shorter time-horizons could be related to more exploitation (see Berger-Tal et al., 2014). Accordingly, individuals whose resources in terms of time are more limited may be more prone to rely strongly on recommendations, leading them to get stuck in inferior options. This also relates well to the finding that transmission chains of individuals who could only try out a more limited number of options were less successful at converging on the global optimum (Yahosseini & Moussaïd, 2020). Thus, individuals' decreased capacity to explore novel information may be an important factor that affects the success of converging to the best solution in large environments.

Beyond that, advisors who depend on the performance of advice recipients may change their search behavior. Thus, holding those accountable who provide social information to others can change their exploration behavior and may change the quality of the recommendations they pass on (Dana & Cain, 2015). In fact, advisors who depend on their recipients may become more likely to pass on information that is also novel to the recipient. This may fuel collective transmission dynamics by providing the system with valuable information.

Future research could harvest the benefits of computational models for our understanding of group-level and collective phenomena (see Goldstone & Gureckis, 2009). Based on the present work, one could simulate data and generate predictions for different time-horizons, and previous experience levels.

In summary, we observe that individuals integrate individual and social information when they search large, multi-modal landscapes. Overall, individuals rely strongly on their own information, echoing tendencies for the underuse of social information (e.g., Efferson et al., 2008; Mesoudi, 2011; Morgan et al., 2012) and egocentric discounting of advice (Yaniv & Kleinberger, 2000). Crucially, improving our understanding of individual-level processes contributes to our understanding of phenomena on the level of larger collectives and groups. Individual search processes are affected by knowledge about interdependence and social information about the environment. This is reminiscent of our findings on coordination under competition, since benefits from social information are limited there as well and interdependence may affect exploration. As a final social context, I will discuss allocation decisions in which other individuals completely depend on a decision maker and social information consists in information about the other individuals.

## 1.6 Allocation of scarce resources

In the previous contexts outcomes were determined jointly, or individuals could choose to what degree they give up control about their outcome by relying on information provided by others. Yet many decisions require unilateral allocation of resources to others. Whereas the role of physicians increasingly shifted towards providing recommendations, they often still decide on behalf of their patients (Gurmankin et al., 2002). During the COVID-19 pandemic, for instance, doctors were repeatedly forced to make decisions to who they allocate limited life-saving resources (e.g., Hick et al., 2020; Solnica et al., 2020; White & Lo, 2020). Similarly, in the debate about self-driving cars the issue has been raised that they may need to decide who to save in case of an accident (Awad et al., 2018). Thus, sometimes individual agents find themselves in a situation, in which we have to make decisions about others.

Decisions about others involve an extreme form of dependence, characterized by pure partner control (Balliet et al., 2017; Kelley & Thibaut, 1978). The outcome of one individual is completely determined by the other. Such extreme forms of interdependence are captured by the dictator game (see Camerer, 2003 and Engel, 2011 for reviews). An individual has a fixed endowment (i.e., a resource) and decides how much another individual receives of this endowment. Think of parents, who allocate their limited time between themselves and their children (see Hertwig & Hoffrage, 2012) or donations to charity, tapping into considerations of fairness and altruism.

Yet, the consequences of allocation decisions may go beyond the allocation of time or money, where one individual will receive less than another or even nothing. Like in the COVID-19 example, allocation of medical resources may determine life or death.

Situations in which individuals decide about others' life or death have been extensively investigated in the context of sacrificial dilemmas (e.g., Everett & Kahane, 2020; Waldmann et al., 2012). Maybe the most prototypical variant of a sacrificial dilemma is the *trolley dilemma* (Foot, 1967; Thomson, 1985). The dramatic scenario typically unfolds as follows: A runaway trolley is heading for a group of 5 people on the tracks, who would be killed. The decision maker is close to a switch, which can divert the trolley to another track, where only a single person would be hit and killed. Thereby, decision makers decide about the fate of multiple individuals.

How should individuals behave? Utilitarian norms prescribe to throw the switch in order to save the larger number of people (Bentham, 1983; Mill, 1863; Singer, 2011 cf. Kahane et al., 2018). Accordingly, individuals should weigh cost and benefits to contribute to the greater good, which in this case simply means saving a larger number of people. Utilitarian norms are typically contrasted with deontological norms, which prescribe or forbid certain acts, like directly harming others.

In the version above, individuals typically throw the switch to save the larger number of individuals. Yet, in the footbridge version of the dilemma individuals have to throw a large man from a footbridge to stop the trolley threatening 5 other people. Here, a considerably smaller proportion of individuals is willing to save the larger number of people (e.g., Waldmann et al., 2012), even if the numbers of individuals saved and killed are the same. This finding has been interpreted as evidence that individuals follow a deontologic norm to avoid harming others.

Yet, there is likely more to moral reasoning than following utilitarian or deontologic norms. Likely moral decisions are based on intuitions (Cushman et al., 2006; Haidt, 2001) and emotions (Greene & Haidt, 2002), a universal moral grammar (Mikhail, 2007), or simple heuristics (Gigerenzer, 2010; Sunstein, 2005). Thereby various factors beyond saving the larger number of people affect moral reasoning (Christensen & Gomila, 2012; Rai & Fiske, 2011; Van Bavel et al., 2015). Among others, factors include the relatedness of individuals to the decision targets or victims. Similarly, individuals are less likely to save others from extreme out-groups (Cikara et al., 2010). Finally, there is indication that hypothetical sacrificial dilemmas are treated differently (Bauman et al., 2014; Bostyn et al., 2018) and that they neglect relevant information about the context (Camerer & Mobbs, 2017; Carnes et al., 2015; FeldmanHall et al., 2012).

The influence of various factors also reflects in allocation decisions beyond sacrificial dilemmas. In allocation decisions like the dictator or ultimatum game individuals apply other fairness norms to individuals who do not belong to their group (Kubota et al., 2013). Even if the decision maker is not affected themselves, they treat similar individuals more favorably (Everett et al., 2015; Tajfel et al., 1971; Yamagishi & Kiyonari, 2000). Thus, individuals deciding about others rely on various sources of information, including individual preferences, as well as situational information or social information about decision targets.

Here we contribute to our understanding how social information is integrated with other sources of information in allocation decisions. The following summarizes results from a repeated cross-sectional study that manipulated social information, the decision situations, and measured features of the decision maker as well as external influences, which is detailed in Chapter 4. Specifically, we investigate allocation of scarce resources in the context of triage decisions during the COVID-19 pandemic.

### 1.6.1 Summary of research paper 3 on decisions about allocations to others

Scarcity of life-saving medical resources during the COVID-19 pandemic made hypothetical trolley-type life-or-death decisions disturbingly real for some doctors (e.g., Solnica et al., 2020). Based on existing utilitarian considerations in this context (Persad et al., 2009), guidelines and recommendations about who should be saved were developed (e.g., Emanuel et al., 2020; Savulescu et al., 2020). Several studies concluded that individuals largely decided in line with guidelines and utilitarian considerations (Fallucchi et al., 2020; Jin et al., 2021; Wilkinson et al., 2020) and that the decision context matters for individuals' judgments what is appropriate (Kneer & Hannikainen, 2021). Yet, it is still unclear, which information actually plays a role for individuals' allocation decisions.

We investigated how individuals allocated scarce artificial ventilation between patient pairs who differed in one focal feature. We gave participants the option to allocate the ventilator randomly to express indifference. Patient pairs differed on one of various features, which were related to their health status (e.g., whether they have diabetes), but also to group membership (e.g., ethnicity), and cooperative or moral behavior (e.g., whether they evade paying taxes). Additionally, we measured individuals' perceptions of the ongoing crisis, as well as characteristics that could be related to some of the patient features (e.g., ethnicity, alcohol consumption, vaccination behavior).

We recruited a large U.S.-based online sample. Our main outcome measure was how likely a given patient would survive given participants' allocation decisions. We find that individuals make use of the patient features in order to make their decision. In line with utilitarian considerations they are more likely to save those individuals who might have a higher chance of survival. Yet, also other considerations appear to enter decisions. Individuals prefer to save patients who were born in the U.S. and those who were described as paying their taxes regularly instead of evading them.

Context information also appears to enter the decision. Interestingly, if individuals are affected themselves, because they had a COVID-19 infection, the probability that the ventilator is allocated to the favored individual is decreased. Similarly, individuals who perceive infections as more severe or more likely, they less likely allocate the ventilator to the favored individual. At the same time, individuals who perceive the crisis as more severe are more likely to save a health professional than a police officer. Thus, information from the participants real-world context appear to be integrated into allocation decisions.

Moreover, we find evidence for some features that the probability that a patient is saved is increased among individuals who share the feature or consider a the feature relevant. For example, among conservatives individuals born in the U.S. are saved more likely.

Likewise, among those who approve more of flu vaccination an individual vaccinated against the flu survives with higher probability. Finally, we find that the preference for saving certain patients also reflects in decisions about withdrawing the ventilator.

Thereby, it is not entirely clear whether the effects reflect in-group favoritism or punishment of individuals, who show non-cooperative behaviors like tax evasion or fail to show cooperative behavior like flu vaccination. It likely is a combination of both. Irrespective of the answer to this question, results shows that individuals integrate social information about the targets with situational information and private information, including their own perceptions and characteristics.

### 1.6.2 Implications and future directions

Also in decisions in which others unilaterally depend on the decision maker, individuals integrate social information with private information, like their own perceptions of the situation or the features of the social others. This corroborates and extends findings from other studies of allocation decisions in the context of COVID-19, showing that individuals prefer others who share their nationality (Jin et al., 2021). Similarly, it reflects findings that individuals' group membership enters decisions in sacrificial dilemmas (Cikara et al., 2010). The findings also relate well to other allocation games, where individuals are more likely allocate resources to others with whom they share even only minimal similarities like taste of art (Tajfel et al., 1971; Yamagishi & Kiyonari, 2000). In a similar vein, individuals change their allocations in the dictator game as a function of how deserving a recipient is (Engel, 2011). We contribute to these findings and show that individuals integrate various sources of information to reach their allocation decision.

It is crucial to raise the question whether individuals actually perceive themselves as independent from their decision target. It seems apparent, that the decisions do not entail interdependence of decision makers' and patients' outcomes. However, decision makers may not be, or at least not perceive, themselves as independent and unaffected from their decision. Obviously, hypothetical decisions cannot actually affect the decision makers. Yet, the real world counterparts of these decisions would entail at least indirect forms of interdependence.

First, in real interactions interdependence can emerge through additional processes. Repeated interactions and direct reciprocity, as well as benefits for ones' group affect decisions towards others (Everett et al., 2015). Individuals are more likely to interact with others with whom they have a shared identity. Thus, according to norms of reciprocity, it may be more important to cooperate with (e.g., save) an individual of one's in-group. Next, saving similar individuals can benefit oneself indirectly through benefiting one's group. If individuals preferentially allocate resources to individuals of their in-group, the group as a whole could be more successful (see Balliet et al., 2014). In a related vein, integration of social information in the form of others' features can help to uphold cooperation within a group. Punishing norm-violating behavior (e.g., tax evasion) through non-allocation increases resources among individuals who show cooperation. Thus, allocation to norm-conforming individuals can be seen as a form or altruistic (Fehr & Gächter, 2002) or third-party punishment (Fehr & Fischbacher, 2004). Beyond these more subtle forms of potential mutual dependence or interdependence, more directly, advisors could be held accountable for the information they pass on (see Dana & Cain, 2015) or doctors may be held accountable for their decisions. In sum, interdependence in seemingly unilateral allocation decisions could emerge because the decision maker of the response of similar others or benefits to the group. Hence, decision makers may consider interdependence with similar others larger than with dissimilar others.

Overall, it is unlikely that individuals directly consider all of these forms of interdependence in their judgments or decisions. Rather this serves as a pointer to which kinds of dimensions the process underlying allocation decisions could consider. In our study effects of group-membership could both reflect in-group favoritism (Tajfel et al., 1971; Yamagishi & Kiyonari, 2000) or increased third-party punishment (Fehr & Fischbacher, 2004). Surely, it could as well be a combination of both. Therefore, future work should shed light on the underlying processes of the strategies by which individuals integrate considerations like in-group favoritism and punishment of non-cooperators.

Importantly, the present study did not aim to make precise predictions about underlying processes. Yet, there are possible candidates that future work could consider. It has been suggested that moral heuristics (Sunstein, 2005) may simply be social heuristics (Gigerenzer, 2010). For instance, fairness can be easily achieved with the $1/N$ strategy, which allocates the same amount to all available parties. Yet, such a simple heuristic seem not at play in moral decisions in sacrificial dilemmas. Otherwise, individuals would not prefer certain patients over others.

Thus, one should consider processes of information integration. It is plausible that strategies for the integration of information that apply to non-social decisions also apply in the context of social or moral decisions (Tobler et al., 2008). Essentially, allocation decisions about others can be considered as *multi-attribute choice tasks*. Multi-attribute choice tasks typically involve choices between different options which differ on some number of attributes and have been investigated extensively (Gigerenzer & Gaissmaier, 2011; Payne et al., 1988). Whereas the individuals in the above study only differed in one key feature, the use of decision maker characteristics and situational characteristics indicates that additional features entered the final decision. In such tasks, lexical decision strategies like Take-the-Best (TTB) use information in a non-compensatory fashion (Payne et al., 1988). Pieces of information are considered sequentially and a decision for one option is made as soon as the options differ on the first feature. Such simple non-compensatory strategies ignore part of the information (Lee & Cummins, 2004; Newell & Lee, 2011). Thereby, they can be superior to strategies that rely on more information (Gigerenzer & Gaissmaier, 2011).

Individuals may rely on similar strategies to make their allocation decisions. The information based on which decision makers compared patients (options) could include moral norms, like avoiding harm and promoting fairness, as well as patients' survival probability, their norm violations, or societal relevance. As soon as strong information, like avoiding both individuals' deaths, does not apply, individuals likely turn to other cues to make their decision. The development and testing of such process models for information integration in allocation decisions and moral reasoning can yield important insights into moral reasoning and its commonalities with other processes of reasoning and decision making.

What is to be gained from understanding the integration processes underlying moral decision making? In fact, one might think that understanding the rules people use in moral reasoning and allocation decisions could inspire smart design and help to develop valuable decision tools. The decision makers that benefit from these rules need not even be human, as in the case of self driving cars, which may need to decide who to spare in case of an accident (Awad et al., 2018). Understanding the rules by which humans make the same decision could allow to build algorithms which resolve situations of accidents most similar to what actual people would do. Yet, this approach has several problems. In clinical judgments such paramorphic models (Hoffman, 1960), that are built to reflect human judgment, are less accepted than clinical judgment. Moreover, using human decision makers as a base, may perpetuate biases which we find morally unacceptable (Jaques, 2019). Even if it may not be the best way to tackle contemporary challenges, understanding the processes of moral reasoning and decision making is at least invaluable to improve our understanding of human reasoning in social (and non-social) situations.

## 1.7 General discussion

This dissertation addresses decision making in a social world, in which decision makers' outcomes depend on others or others' outcomes depend on them. Specifically, three projects asked how individuals (1) establish coordination under competition, (2) search information when giving and receiving recommendations, and (3) allocate scarce medical resources. Each single project yields important insights about one specific social situation involving dependence or interdependence between individuals. Beyond that, comparing and connecting these situations to each other yields additional insights that contribute to our understanding how individuals make decisions in a social world. For this reason we now consider recurring themes and issues across the three projects.

### 1.7.1 Navigating a social world: Common challenges and strategies

One unifying feature across investigations of coordination, recommendation, and allocation is that challenging situations arise through some kind of scarcity. Scarcity is what creates the need to coordinate on options with limited resources and to allocate a limited number of choices, as well as to dispense medical resources. In fact, challenges arise, since all situations involve one or more dilemma structures. Essentially, scarcity leads to a social dilemma that hinders coordination (Dawes, 1980), the exploration-exploitation dilemma requires the decision which option to click (Cohen et al., 2007; Mehlhorn et al., 2015), and a sacrificial dilemma necessitates to decide which patient to save (Everett & Kahane, 2020). But what does behavior across these situations have in common?

All three chapters indicate that humans rely on social and non-social information and individual capacities to tackle these challenges. Competition for limited resources in Chapter 2, requires individuals to allocate their choices in a coordinated fashion. Thereby, they also face an exploration-exploitation dilemma, since they may aim to find a course of action that is superior to maximizing, like finding a pattern (Fantino & Esfandiari, 2002; Peterson & Ulelha, 1965). Coordination is especially challenging, because they need to navigate the exploration-exploitation dilemma, learn what the optimal distribution is, and tacitly coordinate on exploiting it. Like in previous studies on probability learning, individuals on average learn to select the more likely of two options with a higher probability (Vulkan, 2000). This indicates that individuals' allocation of choices is affected by their rewards and the probability of being rewarded. Second, coordination benefits from mutual observation. The broad conclusion from this is, that decision makers rely on both their capacity for individual learning and social information, if the latter is available.

Exploration and search in the recommendation situation described in Chapter 3, require individuals to allocate their choices to discover good options in a large environment. They face the exploration-exploitation dilemma, because they might benefit from options they know, yet then could miss out on even better options. Advisors who generate a recommendation use information about their interdependence with the recipient of that recommendation. They change their search behavior if they know that their reward will depend on others so, that they could increase their chances of finding the best possible option. Recipients, in turn, are attracted by recommendations and use their ability to generalize, to avoid bad recommendations. Thus, their exploration behavior is affected by both, previous experience and generalization (Gershman & Daw, 2017), as well as recommendations that others provide. In sum, both advisors and recipients rely on individual and social information to inform their exploration behavior.

How individuals rely on social information when they explore a probability learning task or spatially correlated bandits, reveals that individuals are often reluctant to fully rely on social information. This insight corroborates a large body of research which suggests that individuals rely on others but frequently decide against fully exploiting their experience (Acerbi et al., 2016; McElreath et al., 2005; Mesoudi, 2011). A potential explanation involves the gain in information. Individuals generally aim to gather information about unknown options in their environment (e.g., Wilson et al., 2014), and thus resolve uncertainty (see Mehlhorn et al., 2015). Consequently, as soon as social information is perceived to not provide a sufficient amount of additional insights, individuals may be inclined to continue the exploration of other options on their own. Therefore, individuals may have continued to switch between options in a probability learning task and explored further in spatially correlated bandits, despite social information.

Finally, allocation decisions in Chapter 4 require individuals to make the choice which of two patients will be saved. They face a sacrificial dilemma, because their allocation decision requires them to weigh the lives of two individuals against each other. Allocations are clearly affected by information about the patient, but also by perceptions of the decision maker and the fit of decision makers' characteristics with features of the patient.

But why do perceptions of hypothetical others affect behavior? A possible explanation is perceived interdependence. We saw in Chapter 3, that advisors change their exploration behavior, as soon as they know that they depend on others. Beyond that, Chapter 2 suggests that perceptions of interdependence can affect which options individuals select. When individuals explore a probability learning task, at least some individuals behave differently, when they know about others' presence. In fact, this is similar to findings on informational social influence (Deutsch & Gerard, 1955), according to which the mere information what others do can affect behavior. In a similar vein, individuals may rely on strategies that account for perceived interdependence as well as social norms, even if the situation is purely hypothetical. Their behavior may be adapted to situations where reciprocity and benefits for their in-group (Everett et al., 2015) are important for their own success.

In sum, across seemingly different social situations individuals rely on different sources of information. Sources of information include the results of learning and their own previous experience, others' behavior and their recommendations, as well as perceived or actual dependence and interdependence. This raises two crucial questions. First, whether the way individuals integrate different sources of social and non-social information is adaptive or rational. Second, how they integrate the different kinds of information more precisely.

### 1.7.2 Is the way individuals integrate different sources of information rational?

The question whether people navigate the challenges of the social world in a rational manner has no straightforward answer. Based on models of utility maximization one could easily conclude that people behave irrationally. On average they fall short of reward maximizing exploitation behavior (Chapters 2 and 3) and fail to be impartial in their allocation decisions (Chapter 4), violating one norm of utilitarianism (see Kahane et al., 2018).

However, the level on which one maximizes the outcome or utility can be important for the judgment of rationality. This highlights a challenge for investigating decisions in a social world: It is crucial to be aware of the multiple levels on which behavior can have effects. This echoes claims, that it is essential to consider both the individual and the environment to judge rationality (Simon, 1956) and rationality is therefore *ecological* (Gigerenzer et al., 1999). Importantly, the environment thereby also includes the social environment, which needs to be considered to judge rationality (Hertwig et al., 2012). Given that different benchmarks can provide different answers, it is important to not jump on premature conclusions about individuals' rationality (see also Neth et al., 2016).

Being adaptive on the group-level, the failure to show probability maximizing could reflect a partial adaptation to competitive settings as seen in Chapter 2. More generally, the failure to exploit options yields the benefit that a group is provided with novel, up-to-date information. Although costly on the individual level, Chapter 3 illustrates that novel information can be especially beneficial to discover better options among large numbers of alternatives. In fact, such innovations are one of the main drivers of successful cumulative cultural evolution (Legare & Nielsen, 2015). Moreover, in sacrificial dilemmas, decision makers may balance the goal to punish non-cooperators to maintain group functioning (Fehr & Fischbacher, 2004) and benefit their group through parochial altruism (e.g., Bernhard et al., 2006). Although not necessarily impartial, this kind of behavior is still in line with utilitarian cost-benefit calculations (Conway et al., 2018). These considerations indicate, that the level on which success is measured strongly affects the judgment whether behavior is rational. Apparently irrational decisions on the individual level can be adaptive on the group-level, so that a judgment about individuals' (irr-)rationality would surely be premature.

### 1.7.3 Future directions and processes of integrating social and non-social information

Independently of whether behavior is rational or not it is important to ask how individuals integrate the various sources of information. Our understanding of decisions in a social world can vastly benefit from the integration with other paradigms. Research on multiple cue probability learning (MCPL; e.g., Kruschke & Johansen, 1999) can provide useful guidance on the question how different pieces of information are integrated into judgments. Information integration can be done either in a compensatory fashion, which weighs information to make a decision, or in a non-compensatory fashion, which relies on part of the information and ignores the other (Gigerenzer & Gaissmaier, 2011; Lee & Cummins, 2004; Newell & Lee, 2011; Payne et al., 1988). Using this type of task, it has been shown that social information is processed differently from non-social information (Puskaric et al., 2018). Beyond that, MCPL can be related to multi-armed bandits (Schulz et al., 2018), since in both cases individuals make inferences about multiple options based on different pieces of information.

I propose that conceiving of multi-armed bandits as similar to MCPL can be another fruitful way of integrating research on social learning and judgment and decision making. In the present research both processes have some plausibility. As detailed above, integration of various pieces of information about others could be a plausible process in sacrificial decisions. More generally, individuals could pool information from social and non-social sources (as assumed in Chapter 3) to inform their estimates of option values. Yet, individuals could also decide which information they consider as more reliable and try to make a difference between options based on that. Unfortunately, also the present results do not allow to reliably disentangle the use of compensatory and non-compensatory processes. Yet, in Chapter 3 the results partially support predictions from compensatory strategies which weight recommendations and estimates generalized from previous experience.

One issue concerning the interpretation of results as evidence for individuals' strategies is, that group-level averages in the presented studies can result from the behavior of single individuals who differ in their behavior or processes that work the same across individuals. Relatedly, it has been shown that individuals differ in how much they rely on social information (Molleman et al., 2019). To disentangle how different individuals integrate information to navigate their social world, future work should continue to address the question, how they treat and integrate different sources of information.

The above discussion shows that combining insights from research on judgment and decision making (like the weighting of attribute features) and research on the use of social information is a fruitful avenue for future research. It highlights, that understanding individual-level processes can be important to better understand collective phenomena. In fact, the strategies that individuals use can create phenomena that emerge from individual interaction, as in the case of many types of coordination (Goldstone & Gureckis, 2009; Moussaid et al., 2009). Moving between the levels on which behavior has consequences provides insights into whether and how strategies are adapted to decisions in a social world. Overall, investigating behavior in situations of varying dependence from different angles and from the perspectives of different research traditions will provide a more comprehensive picture of how individuals make decisions in a social world.

# 2

# Coordination under competition

## An ideal match? Learning to distribute choices under undisclosed competition with limited social information

*Joint work with Wolfgang Gaissmaier,[1,2,3,4] and Hansjörg Neth.[3,4]*

[1] *Graduate School of Decision Sciences, Universität Konstanz*
[2] *Graduate School of the Social and Behavioural Sciences, Universität Konstanz*
[3] *Department of Psychology, Universität Konstanz*
[4] *Centre for the Advanced Study of Collective Behaviour, Universität Konstanz*

## Abstract

When competing with others for limited resources one should not necessarily exploit the option that is most likely rewarding. Individuals have been shown to adapt to this challenge, even under risk and uncertainty. Yet, it has not been systematically investigated how individuals learn to adapt to undisclosed competition as a function of information about the social context. We report 2 online experiments in which participants made 200 predictions in a binary probability learning task. We show that individuals adapt to competition by selecting the more likely event less frequently than without competition. Yet, most groups fail to optimally distribute their choices and coordination remains unstable. Benefits from observing others and influences of knowing about others' presence without observing them do not reflect in systematic differences in individual choice proportions. Although, group-level coordination is more likely under mutual observation. Moreover, knowing about others' presence may prevent maximizing strategies and could reflect a generalized adaptation to competition. Finally, individual differences in cognitive reflection could facilitate coordination in the absence of social information but hinder coordination in its presence. Overall, adaptive information use and robust strategies may help to tackle the challenging task of learning to coordinate when competing for limited resources under risk and uncertainty.

## 2.1 Introduction

Competition for limited resources creates the need to coordinate behavior. For instance, it may pay off if some individuals accept the less effective of two vaccines, since they will compete with fewer others and ensure that the available supplies are efficiently used. Yet, the better option or where one is more likely rewarded is often not known and needs to be learned from experience (Hertwig et al., 2004; Hertwig & Erev, 2009). Additionally, interdependence may need to be learned or inferred from social information (see Balliet et al., 2017).

In most previous research on coordination under risk and uncertainty individuals were typically aware of their interdependence and information about interaction partners was available (Goldstone & Ashpole, 2004; Goldstone et al., 2005; Rapoport et al., 2002; Schulze & Newell, 2015; Schulze et al., 2015; Selten et al., 2007). Learning about interdependence without knowing about others' presence has only been investigated in tasks with non-probabilistic outcomes (Kelley et al., 1962). We ask how individuals learn to adapt to situations with probabilistic outcomes, in which undisclosed competition creates reward interdependence.

We address this question using a probability learning task (e.g., Edwards, 1961; Estes, 1976). Individuals make repeated predictions between two mutually exclusive and exhaustive events which occur with different odds (e.g., $p = 2/3$ and $1 - p = 1/3$) and receive equal rewards for correct predictions. This task captures decisions under risk and uncertainty in which reward probabilities have to be learned by trial and error. Individuals could maximize their rewards by *probability maximizing*, that is, exclusively selecting the more likely event. Instead, decision makers often show *probability matching* and select events relative to their reward probability (for a review see Vulkan, 2000). This choice anomaly appears to be highly irrational and has been considered a result of cognitive limitations (Koehler & James, 2009; West & Stanovich, 2003). Yet, it has long been linked to adaptive tendencies (Gaissmaier & Schooler, 2008; Peterson & Ulelha, 1965; Wolford et al., 2004).

Crucially, probability matching has also been linked to adaptive behavior under competition (Gallistel, 1990; Gigerenzer, 1996; Thuijsman et al., 1995). When the task above is played with two others and limited rewards for correct predictions are split, not all individuals fare best by probability maximizing. If two players already select the more likely event, competition decreases its expected reward for a third player, since the expected reward of $\frac{2}{3}r$ is divided by 3 ($\frac{2}{3}r \cdot \frac{1}{3} = \frac{2}{9}r$). This reward is even lower than the expected reward of the less likely option where one would face no competition ($\frac{1}{3}r = \frac{3}{9}r$). This illustrates that probability maximizing is no stable equilibrium, because one individual could benefit from switching to the less likely option (Gallistel, 1990). Moreover, the rewards from the less likely event would be wasted for the group. Thus, probability maximizing does not lead to efficient reward maximization in all settings. Instead of probability maximizing the optimal group-level behavior is *group-matching*.

Group matching results in the *ideal free distribution* (IFD, Fretwell and Lucas, 1969). This theoretical distribution describes that foraging individuals distribute proportional to resources available in each option. The IFD is optimal, because it constitutes an evolutionary stable Nash-equilibrium (Nash, 1951), in which no individual can do better by switching and the resources in the environment are used optimally. In the above task option probabilities correspond to available rewards. Accordingly, the IFD is achieved if 2 individuals select the more likely, and 1 individual the less likely option (because $\frac{2/3}{1/3} = \frac{2}{1}$). Thus, if individuals match the probabilities by stably distributing *themselves* they achieve an optimal state all of the time, $p(IFD) = 1$.

The adaptive value of probability matching is underscored by the finding that large groups of probability matching organisms distribute according to resource quality (Thuijsman et al., 1995). However, if all individual *choices* match the option probabilities this yields an IFD with a probability smaller than 1 (for the above probabilities of $p = 2/3$ and $1 - p = 1/3$ with $p_{IFD} \approx 0.44$). Thus, if payoffs in a probability learning task have to be shared, probability matching is superior to probability maximizing ($P_{IFD} = 0$) but worse than stably distributing according to an IFD.

### 2.1.1 Learning coordination under undisclosed competition

In group foraging tasks, both humans (Kraft & Baum, 2001; Kraft et al., 2002; Sokolowski et al., 1999) and non-human animals (e.g., Kennedy & Gray, 1993) approximate this optimal distribution. It is yet unclear how individuals adapt their behavior if they have to learn about the probabilities of a reward that can be influenced by the behavior of others. Essentially, they play both a game "against nature" and a "social game" (Hertwig & Hoffrage, 2012, p. 27) and face both *environmental uncertainty* about the rewarded option and *strategic uncertainty* about other players' behavior (Meyer et al., 1992; Rapoport et al., 2002). The simplest way of approximating coordination is through always using a win-stay, lose-shift strategy (WSLS), which results in probability matching in probabilistic settings (Thuijsman et al., 1995). Relatedly, it has been argued individuals learn to use WSLS when they are confronted with undisclosed interdependence under certainty (Kelley et al., 1962). This carries the argument to the extreme that probability matching in non-competitive tasks is an overgeneralization from competitive settings (Gigerenzer, 1996; Schulze et al., 2015). Individuals might *always represent* the probability learning situation as competitive and use a simple, yet robust, strategies that always results in probability matching.

Contrary to this strong argument, individuals have been shown to adapt their behavior to competition and are well described by reinforcement learning models (Erev & Rapoport, 1998; Rapoport et al., 2002), which imply that individuals are capable of coordinating based on their reward signal. More recently, dyads in a probability learning task established stable coordination when between group competition provided an incentive to maximize group-level rewards (Schulze & Newell, 2015). In this study, however, coordination typically was costly for one individual. Relatedly, Schulze et al. (2015) show that individuals adaptively avoid sharing with a competitor by strategies that result in probability matching. Apparently, individuals are able to adapt to competition beyond simple, yet robust strategies, and take into account variation in their individual rewards.

Yet, the conditions under which individuals can adapt to competition are unclear. Previous research has neither systematically varied reward interdependence nor the social information available to discover interdependence. Instead, individuals were typically informed about their interdependence and had information about others' behavior at their disposal (Avrahami et al., 2005; Kraft & Baum, 2001; Kraft et al., 2002; Rapoport et al., 2002; Schulze & Newell, 2015; Schulze et al., 2015; Sokolowski et al., 1999). Systematically investigating the role of social information and competition, however, can yield valuable insights into the conditions for successful adaptation to competition, its underlying processes and their adaptation to competition.

### 2.1.2 Which role does social information play for coordination?

Most likely, learning coordination under undisclosed interdependence is affected by information about others' behavior. In multi-armed bandit tasks individuals integrate information provided by others with their own rewards (Efferson et al., 2008; Toyokawa et al., 2019). Thereby, individuals have been shown to rely on social information across various types of games (Molleman et al., 2014). Imitating the behavior of successful others helps individuals to find and exploit the best option (Laland, 2004; Rendell et al., 2010).

Yet, copying others is also related to potentially maladaptive herding (Bikhchandani et al., 1998; Raafat et al., 2009). In competitive situations under risk and uncertainty this can imply that individuals more likely identify and exploit the more likely option through probability maximizing. Relatedly, in non-competitive group decisions with communication the likelihood of an optimal probability maximizing strategy is increased (Schulze et al., 2020; Schulze & Newell, 2016) and successful others may increase the availability of the probability maximizing strategy, which has been shown to increase the number of individuals who learn to maximize (Koehler & James, 2010; Newell et al., 2013). Crucially, if more individuals learned maximizing under competition, this would result in decreased coordination.

On the contrary, coordination could also benefit from information about others' behavior. In group-foraging tasks coordination was facilitated if others could be observed (Goldstone & Ashpole, 2004) and avoidance of large crowds resulted in a stronger use of the smaller of two resources (Goldstone et al., 2005). Yet, when only the resources were visible, an overuse of the larger option was observed (overmatching). In fact, observation of others may be an important prerequisite for stable coordination. Stable behavior eliminates strategic uncertainty and thus, makes the situation more predictable. As soon as one individual stably selects the less likely option, the other two individuals can safely learn probability maximizing. Relatedly, signalling stable behavior helped to establish coordination (Roberts & Goldstone, 2011). Moreover, if others' behavior is stable individuals are even more willing to accept a lower reinforcement (Ochs, 1990) and may thus be more willing to alter their behavior benefiting coordination. Thus, coordination under uncertainty could benefit from avoiding the more likely option or observing others' stable behavior.

Beyond observing others, already the mere knowledge of others' presence could affect coordination behavior. The presence of others is an important condition to infer interdependence (see Balliet et al., 2017). In non-competitive probability learning tasks some individuals may incorrectly assume the presence of others and attribute non-reward to competition, which results in probability matching or at least failure to maximize. In line with this idea, it has been noted that not only real but also imagined others can affect behavior (see Raafat et al., 2009), like in social influence for conformity (Cialdini & Goldstein, 2004). Relatedly, the mere knowledge about competition can also affect the behavior of non-competitors (KC et al., 2018). Information about the presence of others could make a competitive representation more likely and thus lead to behavior that is adapted to competition. Specifically, when aware of others, more individuals should fall short of probability maximizing than in a standard probability matching task. Accordingly, awareness about others should yield more coordination under competition but less maximizing in the absence of competition.

### 2.1.3 The present study

In two experiments we use a probability learning task to investigate how individuals adapt to competition and to what degree they achieve coordination. Therefore, we vary whether the reward structure is competitive or not. In the non-competitive reward structure the available rewards in each option are unlimited, so that a correct prediction will yield the same reward, independently of the other players' choices. In the competitive reward structure available rewards in each option are limited, so that the reward of a correct prediction depends on the other players' choices. As our second experimental factor we vary what individuals know about the other players. Either they are not informed that other players are there (*unaware*-condition), they are only informed about the presence of other players (*aware*-condition) or can observe the other players' behavior and rewards (*observation*-condition).

The unaware condition allows us to address, whether individuals learn to adapt to competition in the complete absence of social cues and whether observation of others is a prerequisite for successful coordination. Comparisons between competitive and non-competitive settings serve as a benchmark for whether something different is learned under competition or whether individuals surprisingly always follow robust WSLS strategies, which ignore reward magnitude and always result in probability matching. Moreover, these comparisons allow us to assess whether over-generalizations from competitive settings affect choice behavior, if cues about the presence of others imply a social situation. Thus, we expect effects of competition and social information, as well as their interactions.

In Experiment 1 one human participant plays together with two computerized learning agents to address the question, whether they adapt their behavior and which role social information plays. To additionally address the question how coordination is achieved, in Experiment 2 humans play in groups of three.

## 2.2 Experiment 1

Experiment 1 sought to clarify three main issues: First, whether individuals adapt their behavior to competitive settings, second, if information about the presence of other players already affects behavior, even in the absence of competition, and third, to what degree individuals benefit from observing others both in the presence and absence of competition.

One human player played together with two computerized agents. The agents' parameters were determined to obtain agents which learned either relatively stable probability maximizing or matching on average. This distribution reflects previous research which suggests that there is about an equal chance of encountering probability maximizers and matches (e.g., Koehler and James, 2010 report less than 50% maximizers and around 25% probability matchers). The use of agents balanced control of the social environment and uncertainty due to learning and allowed to simplify the recruitment process. Individuals in the non-competitive reward setting could typically benefit from observing the maximizing agent. The competitive reward setting involved the challenge that one of the agents did not settle for an option but retained a fair amount of noise.

## 2.3 Methods

### 2.3.1 Participants

We recruited 310 participants on Amazon Mechanical Turk (134 Female, mean age 36.85 years, range 19 to 72 years). We excluded participants who failed a Captcha or incorrectly answered a comprehension check more than 4 times. Participants received a turn-up fee of US-\$1.00 plus an average performance-based bonus between US-\$1.59 and 3.34 ($M$ = 2.27, $SD$ = 0.24). The experiment was conducted with the software oTree (Chen et al., 2016).

### 2.3.2 Design

Participants played 200 trials of a probability learning task. In each trial they clicked one of 2 buttons to predict which of 2 mutually exclusive events will occur and received rewards for correct predictions. Events occurred with a probability of 2/3 and 1/3 respectively. Participants were informed that events differ in their probability of occurrence and that the rewarded event in any round is independent of the previous rounds (examples of the task screen can be found in Figure A.1).

Each participant played in a group, together with two computerized agents. Each of the agents belonged to one of two different types. The two agent types were derived from preliminary simulations of agent behavior in the Projective Simulation model (Briegel and De las Cuevas, 2012; for details on our agents see Appendix A.2 on page 95). In the absence of competition, the first type learned a relatively stable behavior whereas the other type learned probability matching (on average respectively). The agents learned together with the human participant based on the rewards they received. Thus, they modelled interdependence with other learners as an important aspect of social situations.

We investigated learning over time and varied social information and competition in a full-factorial 3(unaware vs. aware vs. observation) × 2(competition vs. no-competition) between-subjects design. In the *unaware* condition participants were not informed about the presence of the other players, whereas in the *aware* condition participants were told, that they play the task with two other players and were assigned a number (see upper panel in Figure A.1). In the *observation* condition they also could observe the other players' choices, rewards and total rewards (see lower panel in Figure A.1). Thereby, participants in the aware and observation condition were informed that they play together with learning, computerized agents.

Moreover, participants either shared the points from correct predictions with all others who made a correct prediction on the same trial (*competition*) or everyone received the full number of points (*no competition*). When points were shared a correct prediction yielded a total of 30 points, whereas if rewards were not shared a correct predictions yielded 15 points. Thus, the expected reward for the respective group (and individually) optimal strategy (forming an IFD under competition or probability maximizing in the absence of competition) was held constant at 10 points per trial. Participants were not informed whether they experienced a competitive or non-competitive reward structure.

### 2.3.3 Procedure

After providing informed consent participants provided basic demographic information and information on sightedness. Then they received task instructions and answered 3 understanding questions, they had to answer correctly before proceeding. After the main task participants answered questions about the strategies they used and how they perceived the computerized agents. After that they responded to the Cognitive Reflection Task (CRT, Frederick, 2005) in a version developed by Mata et al. (2013). Cognitive Reflection has been related to increased probability maximizing in previous studies on probability matching (Koehler & James, 2010). Finally, they could provide comments about the study and were then fully debriefed about the study goals and received information about the condition they were in.

### 2.3.4 Data analyses

Our main dependent variable are participants' choice proportions of the more likely event. We analyzed the data with generalized mixed logistic models in the package `lme4` (Bates et al., 2015) in R (R Core Team, 2018). Logistic models better reflect the binary nature and are better suited for the analysis of proportions than traditional ANOVAs (Jaeger, 2008). We added participant random intercepts and random slopes for round to account for non-interdependence of observations. The predictor for round is centered to set the intercept at 100 rounds and scaled to steps of 50 rounds. Thus, round-related predictors denote the change in the probability of selecting the more likely option in 50 rounds. Competition and social information are entered as effect coded. To test our expectations, we test the independent contributions of the fixed effects of round, competition, social information and their interactions by using Wald-$\chi^2$-tests from the package `car` (Fox & Weisberg, 2019). To probe the interactions, we analyze simple effects and slopes with the package `emmeans` (Lenth, 2019).

## 2.4 Results

Effects for round, $\chi^2(1) = 121.968$, $p < 0.001$, competition, $\chi^2(1) = 5.140$, $p = 0.023$, and their interaction, $\chi^2(1) = 5.826$, $p = 0.016$ improve a generalized logistic mixed model predicting the probability that the participant selects the more likely option (for Model tables see Table A.3 on page 97, models 1–3). We do not find evidence for a main effect of social information, $\chi^2(2) = 4.190$, $p = 0.123$, or its interaction with round, $\chi^2(2) = 2.055$, $p = 0.358$. $b$-values are on the logit scale.

Overall, a positive effect of round indicates that human participants learn to select the more likely option with higher probability over time, $b = 0.446$, $SE = 0.040$, $OR = 1.562$, Wald $Z = 11.044$. If there is competition, individuals select the more likely event with a lower probability, $b = -0.186$, $SE = 0.082$, $OR = 0.831$, Wald $Z = -2.267$, and the interaction of competition and round indicates that choices of the more likely option increase less under competition, $b = -0.096$, $SE = 0.040$, $OR = 0.909$, Wald $Z = -2.414$. This indicates that participants learn less probability maximizing under competition.

Contrary to our expectations we do not find evidence for effects of social information. Nevertheless, additional comparisons suggest that the probability of choosing the more likely option is lower in the aware condition, both if there was no competition, $OR = 0.616$, $SE = 0.173$, Wald $Z = -1.720$, $p = 0.518$, and if there was, $OR = 0.776$, $SE = 0.210$, Wald $Z = -0.937$, $p = 0.937$. However, neither effect can be reliably distinguished from zero. We take this to indicate that most of our participants did neither make substantive use of learning through copying, nor did they select the more likely option less frequently as a function of being aware of others' presence.

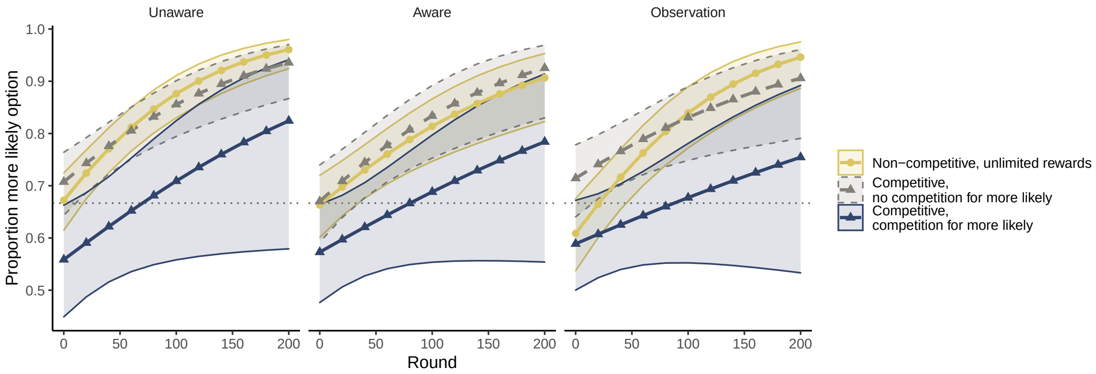

Figure 2.1: Proportion of more likely option choices as a function of rounds, available social information and competition. The competition condition is additionally split by whether both computerized agents were more likely to select the more likely option (dashed line). Only if this necessary prerequisite is fulfilled, we observe a difference between competition and settings with unlimited rewards. The dotted horizontal reference line indicates the probability expected from probability matching. Ribbons correspond to 95% confidence intervals.

**Adaptiveness of selecting the less likely option** It was unclear whether participants' decreased maximizing behavior originated from an adaptive tendency or reflected noisy behavior. Based on their rewards, the computerized agents learned either to select the more or the less likely option with a higher probability. A necessary (yet not sufficient) condition for optimal avoidance of maximizing is that both other players select the more likely option above chance. Based on agents' choice proportions across rounds we classified *post hoc* whether the condition was met.

This distinction was used to modify our experimental competition condition to whether competition for the more likely option was plausible or not (note: without competition maximizing is unconditionally optimal). A model with this modified predictor provided a better fit to the data than a model with our original manipulation only, $\chi^2(6) = 19.961$, $p = 0.003$, $\Delta AIC = 7.96$. This indicator, $\chi^2(2) = 19.942$, $p < 0.001$, as well as its interaction with round, $\chi^2(2) = 7.847$, $p = 0.020$, contribute independently to model fit (see Table A.3 on page 97, model 4).

Figure 2.1 illustrates the results. Only when both agents select the more likely option with a probability greater than chance participants select the more likely option less likely than on average, $b = -0.887$, $SE = 0.210$, $OR = 0.412$, Wald $Z = -4.231$, $p < 0.001$, and learn to probability maximize less over time, $b = -0.286$, $SE = 0.103$, $OR = 0.751$, Wald $Z = -2.765$, $p = 0.006$. In the competitive setting we observe no difference to the average probability of selecting the more likely option; neither overall, $b = -0.034$, $SE = 0.183$, $OR = 0.966$, Wald $Z = -0.187$, $p = 0.852$, nor with regard to the interaction with time, $b = -0.129$, $SE = 0.091$, $OR = 0.879$, Wald $Z = -1.423$, $p = 0.155$.

Of course, this approach is not without problems, because the distinction was introduced *post hoc*. Hence, it is unclear whether the agents' behavior was responsible for participants' behavior or whether participants led the learning agents to avoid the more likely option through their behavior. Nevertheless, those participants who select the more likely option less frequently do so in situations in which this can be adaptive. The absence of an effect of social information, $\chi^2(2) = 3.157$, $p = 0.206$, (and its higher-level interactions), indicates that the change in participants' behavior is due to their individual rewards.

**The role of CRT** Finally, we also explored how individual differences in Cognitive Reflection affect choice behavior. Previous research indicates that higher levels of cognitive reflection are related to more probability maximizing (Koehler & James, 2010). In line with that, a main effect of participants' CRT-score, $\chi^2(1) = 8.663$, $p = 0.003$, indicates that a larger number of correctly answered items is related to more choices of the more likely option, $b = 0.205$, $SE = 0.070$, $OR = 1.228$, Wald $Z = 2.943$ (see Table A.3 on page 97, model 5).

## 2.5 Discussion

As expected, we find evidence that individuals adjust their behavior to competitive settings by selecting the more likely option less frequently. Thereby, behavior is sensitive to the agents' choice behavior. Yet, we rarely observe stable individual minimizing behavior (the highest probability of more likely option choices below chance is at about 16% in the competitive setting if others could be observed). One natural explanation for this is that one of the agents typically retained a large level of noise in its behavior. Thus, it was impossible to achieve a stable distribution.

Surprisingly, we do not find systematic effects of the available social information. A possible explanation is that individuals did not consider the computerized agents as social context. Moreover, coordination may rather benefit from mutual observation, since individuals may signal their willingness to commit to one option (Roberts & Goldstone, 2011). Yet, the computerized agents did not have access to the behavior of human participants and learned only based on their own rewards.

Accordingly, the present results do not allow us to directly investigate whether and how groups of humans can achieve stable coordination. To address this issue we conducted a study with the same design investigating groups of humans. This also allows us to rule out a different treatment of computerized agents as an explanation for our results.

## 2.6 Experiment 2

Experiment 2 extended Experiment 1 to groups of three human players. Beyond the question whether individuals adapt their behavior to competitive settings and if their individual-level behavior is affected by social information, it allows us to address two additional questions. The first question is whether groups manage to coordinate because individuals distribute their choices or rather because they distribute themselves. The second question is, if the absent effect of social information can be explained by the use of computerized agents in Experiment 1.

We expected to replicate the effect that adaptation to competitive settings will reflect in an increasing proportion of group matching and decreased choices of the more likely option under competition. On the group level we expected the proportion of group matching under competition to be larger than the proportion to be expected if all individuals choose between the two options randomly, $p(IFD|random) = 0.375$. If individuals don't neglect social information from other human players in the absence of competition, they should show increased maximizing when able to observe others. Also under competition we expected that group-level coordination benefits from observing others because it allows to display stability. These predictions were preregistered (https://osf.io/sr795). Beyond that, we address how individual differences in Cognitive Reflection affect coordination. The CRT has been implied to affect individual maximizing behavior across conditions in Experiment 1. It is both plausible that increased probability maximizing among those high in Cognitive Reflection (Koehler & James, 2010) hinders successful coordination, since groups of maximizers never achieve an optimal distribution. Alternatively, it could help to coordinate successfully because increased maximizing under competition can increase the expected reward of selecting the less likely option.

## 2.7 Methods

We used the same choice task and design as in Experiment 1 (example screens in Figure A.1). The crucial difference was that players now played in groups of 3 humans, instead of playing with learning computerized agents.

### 2.7.1 Participants

Originally, we aimed to recruit 540 participants in 180 groups (30 groups of 3 in each of 6 conditions) to be able to detect medium-sized effects in an ANOVA-design. We exceeded our goal due to imbalance in group size through randomization error. We recruited 651 participants in 217 groups on Amazon Mechanical Turk (302 Female, mean age 36.3 years, range 18 to 73 years). Participants received a turn-up fee of US-\$1.00 plus an average performance-based bonus between US-\$1.88 and 3.98 ($M$ = 2.76, $SD$ = 0.36). The experiment was conducted with oTree (Chen et al., 2016). We excluded participants who failed a captcha or incorrectly answered a comprehension check more than 4 times. In line with our preregistration we also excluded groups with at least 1 participant who reported poor data quality or spent very short time on the instructions. This resulted in a total sample of 206 groups with 3 participants each (618 participants). We report, where the full sample and this filtered sample diverge.

### 2.7.2 Procedure

The social information condition was determined randomly before each data collection session. In order to be able to group participants without their awareness, participants rated pictures unrelated to the present study until they could be grouped. Participants, and thus their waiting time, were rewarded for each rated picture. As soon as 3 participants were available at the same time, they were grouped, and this group was randomly assigned a competition condition. Participants who could not be grouped were dropped from the study and received their participation fee plus any rewards accumulated in the picture rating task. We do not find an effect of participants' waiting time on their choice proportions, $r = 0.004$, $p = 0.924$.

After the grouping procedure, and prior to the main task participants were shown a summary of task instructions. In the main task, participants made 200 predictions which of 2 mutually exclusive events will occur. A 6-second time-limit ensured that the times between trials were approximately equal across social information conditions. If participants timed out more than 5 times, indicating non-participation, the group ended and all participants were debriefed and those who did not end their participation received the rewards accumulated so far. Otherwise, the procedure was mostly identical to that of Experiment 1.

Analogously to Experiment 1, after the main task participants answered questions about the strategies they used and how they perceived the other players. Participants also responded to the Cognitive Reflection Task (Frederick, 2005) in the version developed by Mata et al. (2013). Finally, they could provide comments about the study and were then fully debriefed about the study goals and received information about the condition they were in.

### 2.7.3 Data analyses

Our main dependent variable is group-level coordination, defined as reaching the optimal state of an IFD in a given round. Moreover, we report results on individual-level choice proportions of the more likely event. All analyses were conducted in in R (R Core Team, 2018). Deviating from our preregistration, we analyzed the data with generalized mixed logistic models in the package `lme4` (Bates et al., 2015), instead of repeated measures ANOVAs. Logistic models better reflect the binary nature and are better suited for the analysis of proportions than traditional ANOVAs (Jaeger, 2008). Moreover, mixed models allow to account for interdependence between individuals in groups (Gelman & Hill, 2007). We accounted for non-independence between observations by including random effects. For analyses of group-level coordination we added random intercepts and random slopes for round on the group level. For analyses of individual-level choice proportions we entered intercepts and random slopes for round for participants nested in groups. The predictor for round is centered to set the intercept at 100 rounds and scaled to steps of 50 rounds. Thus, round-related predictors denote the estimated change in the probability of forming an IFD or selecting the more likely option in 50 rounds. Competition and social information are entered as effect coded.

We test the independent contributions of the fixed effects of round, competition, social information and their interactions by using Wald-$\chi^2$-tests from the `car`-package (Fox & Weisberg, 2019). The predictor for round is centered to set the intercept at 100 rounds and scaled to steps of 50 rounds so that round-related predictors denote the change in probability in 50 rounds. Competition and social information are entered as effect coded. To probe the interactions of our experimental factors, we additionally use contrasts and post-hoc tests of simple effects and slopes conducted with the package `emmeans` (Lenth, 2019). If one individual timed out in a particular round the complete group was not considered for this particular round.

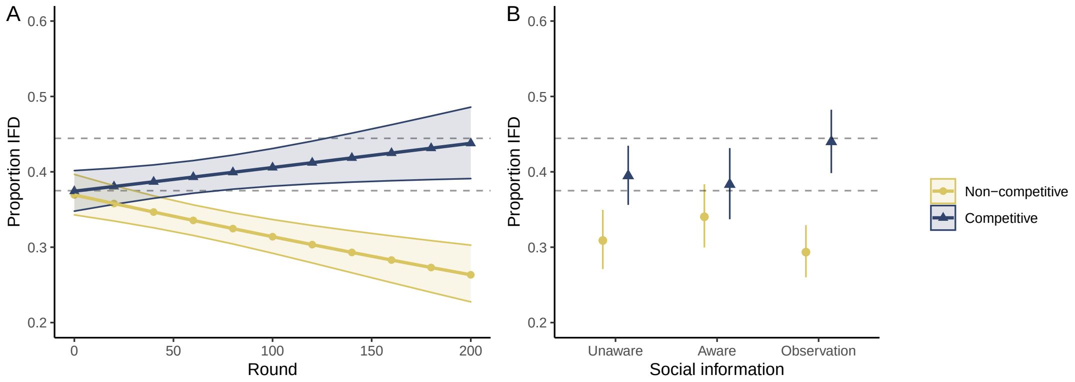

Figure 2.2: Predicted group-level probabilities of reaching a coordinated 2:1-distribution (IFD) as (A) a function of round and competition condition and (B) the interplay of social information and competition. Under competition groups achieve the coordinated distribution more frequently over time. In the unaware and observation condition there is a difference between competitive and non-competitive settings that is absent in the aware-condition. Only in the observation condition groups are more likely to achieve coordination than a group of randomly choosing individuals (lower dashed line). Results are averaged across the variables not displayed. Ribbons and error bars correspond to 95% confidence intervals.

## 2.8 Results

Competition for limited resources and social information reflect in the probability of forming a group-optimal configuration. Under competition, distributing 2:1 to the more and less likely option respectively (IFD) is an optimal (Nash-equilibrium) solution on both the individual and the group level. Whereas probability maximizing is optimal in non-competitive settings but results in a decreased probability of forming an IFD ($p = 0$ for strict maximizing).

We fitted a generalized binary logistic mixed model predicting the probability that a group forms an IFD in a given round (for model tables see Table A.6 on page 100). We find no evidence for main effects of round, $\chi^2(1) = 1.731$, $p = 0.188$, or available social information, $\chi^2(2) = 0.463$, $p = 0.793$. However, we find evidence for effects of competition, $\chi^2(1) = 28.607$, $p < 0.001$, its interaction with round, $\chi^2(1) = 18.821$, $p < 0.001$, and its interaction with available social information, $\chi^2(2) = 6.085$, $p = 0.048$. The latter effect is not entirely robust to inclusion of low-effort participants, $\chi^2(2) = 5.084$, $p = 0.079$. Yet, the general pattern of results is unaffected.

Further tests of the interactions reveal, that optimal IFD-configurations are reached more likely in competitive settings, $b = 0.07$, $CI_{95\%} = [0.01; 0.13]$, $OR = 1.07$, and less likely in non-competitive settings, $b = -0.12$, $CI_{95\%} = [-0.18; -0.06]$, $OR = 0.88$, see Figure 2.2, panel A.

The interaction between available social information and competition reflects that the difference in the probability of forming an IFD varies with the available social information, see Figure 2.2, panel B. Contrary to the expectation that awareness of others increases the probability of achieving an IFD in the absence of competition, the awareness condition does not differ systematically from the unaware condition, $b = -0.14$, $CI_{95\%} = [-0.41; 0.12]$, $OR = 0.87$. Yet, we find evidence for differences between competition and no competition in the unaware, $b = 0.38$, $CI_{95\%} = [0.13; 0.62]$, $OR = 1.46$, and observation condition, $b = 0.64$, $CI_{95\%} = [0.40; 0.88]$, $OR = 1.89$, and no evidence for a systematic difference in the awareness condition, $b = 0.19$, $CI_{95\%} = [-0.09; 0.46]$, $OR = 1.20$. These results could indicate that participants, who are aware of each other, behave similar to competitive situations. This is related to our expectation, that awareness of others makes a competitive representation more likely.

Moreover, only if individuals can observe each other under competition, groups form an IFD more likely than if every individual selected options randomly, see lower dashed line at $p(IFD) = 0.375$ in Figure 2.2, panel B, $b = 0.36$, $CI_{95\%} = [0.08; 0.64]$, $OR = 1.43$. Differences to the expectation from randomly behaving groups are non-systematic in the unawareness, $b = 0.14$, $CI_{95\%} = [-0.12; 0.41]$, $OR = 1.15$, and awareness, $b = 0.18$, $CI_{95\%} = [-0.14; 0.51]$, $OR = 1.20$, conditions. The pre-registered $t$-test leads to the same conclusion. Thus, results indicate that group coordination under competition benefits from the observation of others.

### 2.8.1 Selecting the more likely option

Next, we address how individual-level choice proportions are affected by our manipulations (for model tables see Table A.7 on page 101). Again, effects for round, $\chi^2(1) = 260.495$, $p < 0.001$, competition, $\chi^2(1) = 12.489$, $p < 0.001$, and their interaction, $\chi^2(1) = 13.769$, $p < 0.001$, contribute to the fit of a generalized logistic mixed model with random intercepts and slopes for participants nested on the group level, predicting the probability that a participant selects the more likely option. We do not find evidence for a main effect of social information, $\chi^2(2) = 0.976$, $p = 0.614$, its interaction with round, $\chi^2(2) = 1.092$, $p = 0.579$, or competition, $\chi^2(2) = 1.468$, $p = 0.480$.

Like in Experiment 1, human participants learn to select the more likely option with higher probability over time, $b = 0.576$, $SE = 0.036$, $OR = 1.779$, Wald $Z = 16.140$. Competition lowers the probability that individuals select the more likely event, $b = -0.259$, $SE = 0.073$, $OR = 0.772$, Wald $Z = -3.534$, and the interaction of competition and round indicates that learning is slowed down when rewards are limited, $b = -0.129$, $SE = 0.035$, $OR = 0.879$, Wald $Z = -3.711$. This indicates that participants on average learn less probability maximizing under competition. The absent evidence for effects of social information suggest that individual-level learning and choice behavior are not systematically affected by knowledge about others' presence or observing their behavior.

### 2.8.2 Stability of coordination

Despite evidence that the available social information affects group-level coordination under competition, we did not find evidence for effects on individual-level choice proportions. A potential explanation could be that establishing stable coordination is difficult. Not a single group enters a situation in which 1 individual selected the less likely option with a probability of at least 90% and 2 individuals selected the more likely option with a probability of at least 90%. Yet, under competition 40 of 113 groups (35.4%) establish stochastic coordination. That is, they include 1 individual who selects the less likely option more frequently than chance in the last 50 rounds. Thus, although individuals alter

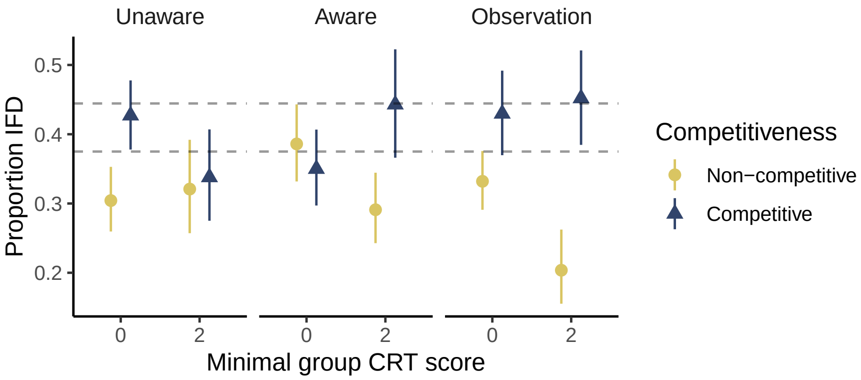

Figure 2.3: Predicted probabilities of reaching an IFD as a function of available social information, competition and lowest CRT-score in the group. CRT scores of 0 and 2 were selected to increase readability. The lower dashed line corresponds to the expected coordination of a group of randomly behaving individuals, the upper dashed line to a group of probability matching individuals. In the unaware condition (leftmost panel) we observe more coordination in competitive than in non-competitive settings among groups with at least one low CRT individual. They reach the coordinated state more frequently than a randomly behaving group. In the awareness condition (middle panel) groups with low CRT individuals coordinate around chance independent of competitiveness. Groups with all individuals above a CRT score of 2 coordinate more under competition when aware of others or observing them. Under observation (right panel) both kinds of groups coordinate more under competition than in its absence.

their behavior in competitive situations and distribute stochastically, we do not observe stable coordination. Rare consistent choices of the less likely option could explain why we do not observe substantial effects on individual-level choice proportions of the more likely option.

### 2.8.3 Cognitive Reflection and coordination

Previous research suggested that differences in CRT are related to less probability maximizing (Koehler & James, 2010). Experiment 1 corroborated this finding for the present task. In line with this idea, an individual-level model predicting choice proportions of the more likely option is improved by including individuals' CRT, $\chi^2(1) = 18.979$, $p < 0.001$, and its interaction with round, $\chi^2(1) = 5.695$, $p = 0.017$, but not by its interaction with competition, $\chi^2(1) = 0.043$, $p = 0.835$ (see Table A.7 on page 101, model 5). Accordingly, individuals with higher CRT-levels select the more likely event more frequently, $b = 0.273$, $SE = 0.063$, $OR = 1.313$, Wald $Z = 4.357$, and learn to do so more frequently over time, $b = 0.070$, $SE = 0.029$, $OR = 1.073$, Wald $Z = 2.386$.

Additional analyses suggest that group-level results are modulated by individual differences in cognitive reflection. We conjectured, that the individual with the lowest CRT could affect coordination because they learn less maximizing and less stable behavior. Therefore, we entered a group's lowest CRT score (group-minimal CRT) to predict how likely the group reaches an IFD. This improves model fit above a model that includes (group average) demographics, $\chi^2(12) = 30.636$, $p = 0.002$, $\Delta AIC = 6.64$ (see Table A.6 on page 100, model 5). We find evidence for a marginal main effect of the number of

CRT-items answered correctly, $\chi^2(1) = 3.512$, $p = 0.061$, its interaction with competition, $\chi^2(1) = 5.326$, $p = 0.021$, and its 3-way interaction with social information and competition, $\chi^2(2) = 13.222$, $p = 0.001$.

Figure 2.3 illustrates how group-minimal CRT modulates the interplay of social information and competition. The following results are corrected for multiple tests by using the Tukey method. Maybe most surprisingly, under unawareness we do not observe a difference between competitive and non-competitive settings if group-minimal CRT is *high*, $OR = 1.08$, $CI_{95\%} = [0.53; 2.21]$, whereas we do if group-minimal CRT is *low*, $OR = 1.70$, $CI_{95\%} = [1.03; 2.82]$. This could be due to more probability maximizing among high-CRT individuals that is not corrected for by choices for the less likely option. Accordingly, low-CRT individuals who are generally less likely than high-CRT individuals to learn maximizing could benefit tacit group coordination in the absence of any social information.

Under awareness, the probability of forming an IFD in competitive and non-competitive settings does not differ systematically if group-minimal CRT is *low*, $OR = 0.86$, $CI_{95\%} = [0.49; 1.50]$. Yet if the group-minimal CRT is large (at least 2 correct items), coordination is about twice as likely under competition than in its absence (yet, also non-systematic), $OR = 1.94$, $CI_{95\%} = [0.99; 3.80]$. Under observation the difference is positive, both among groups with a low, $OR = 1.52$, $CI_{95\%} = [0.90; 2.56]$, and high minimal CRT, $OR = 3.22$, $CI_{95\%} = [1.58; 6.60]$. Since the difference among low minimal CRT groups is non-systematic only due to corrections for multiple tests, $OR = 1.52$, $CI_{95\%} = [1.11; 2.08]$, we cannot conclusively argue for a general disadvantage among low minimal CRT groups. Thus, the competition effect is sensitive to group-minimal CRT under awareness but less clearly under observation. Low CRT individuals could hinder both coordination and maximization under awareness because they fail to learn stable behavior. Yet, if others can be observed, more individuals may be able to override or respond to a competitive representation. Although intriguing, these results are exploratory and should be replicated in a larger sample.

## 2.9 General Discussion

We investigated how undisclosed competition for limited resources and available information on the social context affect learning under risk and uncertainty. In a probability learning task individuals made repeated predictions about two events occurring with different odds. Whereas individuals should optimally learn selecting the more likely option in the absence of competition, under limited resources individuals should optimally distribute themselves in an IFD, or, at least, ensure some selection of the less likely event. In 2 experiments, we find that individuals select the more likely option less frequently under competition than in the absence of competition. Yet, stable coordination is scarce and appears to be challenging. Individuals rather seem to distribute their choices between options rather than distributing themselves. We do not find the expected differences between different levels of available social information. Nevertheless, evidence for variation of the competition effect within social information conditions suggests that information about others' presence and behavior does affect coordination. How individuals adapt their behavior and groups learn to coordinate adds to our understanding of how individuals learn to coordinate in challenging environments under risk and uncertainty where interdependence has to be inferred and options have to be learned from experience.

### 2.9.1 How do people adapt their behavior to competition?

A key explanation for why the adaptation to competition is incomplete, is that this setting is very challenging for various reasons. When coordinating under risk and uncertainty, individuals are not only ignorant about whether an option will be rewarded on the next timestep, but also need to anticipate other players' next move on top of that. Essentially, they play both a game against nature and a strategic game (Hertwig & Hoffrage, 2012) and face environmental and strategic uncertainty (Meyer et al., 1992; Rapoport et al., 2002). Similarly, larger groups have difficulty in establishing coordination if they need to learn navigating an environment in which nature's next move cannot reliably be predicted (Kraft et al., 2002). However, dyads have been shown to establish coordination in a probability learning task (Schulze & Newell, 2015). Why do individuals in our task still fail to do so, even if they can observe others?

First, participants were not informed about their interdependence, unlike in Schulze et al. (2015) and other coordination tasks (e.g., Goldstone & Ashpole, 2004; Goldstone et al., 2005; Kraft & Baum, 2001; Kraft et al., 2002; Rapoport et al., 2002; Selten et al., 2007; Sokolowski et al., 1999). To coordinate stably, 2 individuals have to simultaneously learn maximizing, while 1 individual learns to select the less likely option. One way of achieving this is that individuals infer interdependence and learn what the best option is and tacitly distribute themselves. Yet, even if individuals could infer interdependence through observing others, individuals seemed to have difficulty to establish stable coordination.

Moreover, inferring interdependence is likely challenging, because it essentially requires to disentangle environmental and strategic uncertainty. Specifically, individuals must connect changes in reward magnitude to other players' behavior, whereas they must attribute non-rewarding rounds to uncertainty in the environment. Without information about others this is even more challenging, because individuals cannot directly learn about others' behavior. Instead they need to rely on their own rewards and their situational representation to infer competition and adapt their own behavior accordingly.

Even if others can be observed, persistent (strategic) uncertainty about others' behavior may prevent individuals from achieving stable coordination. Dyads in Schulze and Newell (2015) only established stable coordination if they had an incentive to maximize group rewards and tended towards probability maximizing otherwise. Similarly, in the present task individuals should have had an incentive to coordinate in a stable IFD, because it was both the individually and group-optimal distribution. Yet, as long as the other individuals do not reliably select the more likely option, selection of the less likely option comes at a cost for ones' rewards. More generally, choosing only the less likely option may be also perceived as individually costly, because it means to forego a higher probability of receiving any reward. This could be related to the persistent tendency to weight the more likely outcome more strongly in decisions from experience (Hertwig et al., 2004). Moreover, individuals were shown to accept lower reinforcements if others behaved stable and were to be trusted (Ochs, 1990). An anonymous online setting without communication made it difficult to fulfill the condition of trust. Thus, low levels of coordination could additionally reflect cooperation failure. Similar to the phenomenon of melioration (Herrnstein & Prelec, 1991) individuals may fail to forego more frequent immediate rewards in favor of less frequent but higher sustained rewards when coordination is eventually established.

Another explanation why individuals fail to establish stable coordination is that individuals do not achieve stable behavior in the first place. Signalling stable behavior has been shown to facilitate cooperative outcomes (Roberts & Goldstone, 2011). Yet, it has been shown that higher payoff variability in probability learning is related to slower learning and more random behavior (see Erev & Barron, 2005). Competition can induce additional variability in payoffs through splitting a payoff with others like in the present task. Moreover, in other competitive situations retaining a certain degree of randomness can be adaptive, for instance to avoid predators (Avrahami et al., 2005). For these reasons individuals may have failed to behave predictably and thus hindered successful coordination.

Finally, learning to coordinate within 200 rounds posed an additional challenge, because information is limited and successful states may have been encountered too rarely to allow learning both interdependence and the structure of the environment. Relatedly, more trials can make probability matching disappear (Edwards, 1961). In a similar vein, learning stable behavior and establishing coordination may have taken too long to pay off in the present number of trials (Sims et al., 2013). In the light of these challenges it may actually seem surprising that individuals learned anything although competition was not disclosed. Nevertheless, some groups even figured out some level of divergence, which indicates that establishing coordination is generally possible. One reason may be strategies that even function without observation, another the adaptive use of social information.

But what might individuals do to tackle the challenge of adapting to undisclosed competition under risk and uncertainty? It has been argued that situations of uncertainty require simple heuristics (Gigerenzer, 2020; Volz & Gigerenzer, 2012). Learning clearly suggests that simple WSLS strategies are at least not broadly used from the start. Yet, individuals may learn such strategies over the course of the game (Kelley et al., 1962). Alternatively, basic reinforcement learning processes may operate, as well as regret-based processes which respond to earning higher reinforcement in the unselected option (Kareev et al., 2019). Future investigation should aim to disentangle the role of this kind of processes for learning competition under risk and uncertainty.

### 2.9.2 How does social information affect undisclosed coordination?

The effects of social information and their absence, allow further inferences about which information individuals rely on. First, individuals who can observe the other players do not decrease their choices of the more likely option more, than individuals who cannot observe others. Thus, individuals appear to rely little on social information. In fact, individuals in our task behave similarly to individuals who did not receive social information in other tasks. Individuals who could not observe each other in group foraging tasks and had to rely on individual reinforcement also overused the more abundant resource (Goldstone et al., 2005). Clearly, we also do not observe increased maximizing under competition, which might indicate availability of the maximizing strategy through social information (Newell et al., 2013; Schulze et al., 2020; Schulze & Newell, 2016). Accordingly, decreased choices of the more likely event can largely be attributed to individual reinforcement learning or robust strategies like WSLS.

Alternatively, the time-constraint in Experiment 2 may explain the underuse of social information. Yet, in Experiment 1 there was no time-constraint and still individuals failed to benefit from copying the maximizing agent in the absence of competition. This failure to adaptively use social information is in line with previous findings. Individuals frequently weight their individual information more strongly, even if this means forgoing potential benefits (Acerbi et al., 2016; Efferson et al., 2008; McElreath et al., 2005; Mesoudi, 2011; Mesoudi & O'Brien, 2008a). Additionally, individuals differ in how much they rely on social information (Molleman et al., 2014). In the light of previous findings our, results likely indicate genuine underweighting of social information among a majority of individuals.

The relative under-utilization of information about others' behavior by some individuals, could be an additional reason why groups rarely achieve stable coordination. Nevertheless, it would be wrong to conclude that social information is neglected entirely. Only with mutual observation, individuals achieve coordination more likely than a group of individuals who make their predictions randomly in Experiment 2. The absence of any systematic effects of observation in Experiment 1 indicates that unilateral observation may not be sufficient. Thus, observational information likely is integrated into learning and decision making by at least some individuals. Most basically, the absence of stable coordination suggests that they may try to avoid the option selected by more others (Goldstone et al., 2005; Schulze et al., 2015); either across multiple rounds or, myopically, in the previous round (Kareev et al., 2019).

Beyond learning from observational information, the mere information about others' presence affects learning and behavior. In the aware-condition, the proportion of optimal group-level distributions does not differ between competitive and non-competitive settings. This could indicate that behavior is based on a default representation, which is competitive. In line with this idea the difference is only absent among groups in which the minimal Cognitive Reflection is low. The CRT is supposed to measure the ability to override an initial response (Frederick, 2005; Toplak et al., 2014). Relatedly, only individuals with higher levels of CRT may be willing or able to override their initial, competitive representation. This aligns with the idea that probability matching is the product of an intuitive response (Koehler & James, 2009).

Finally, the surprising finding that individuals who are low in cognitive reflection facilitate group-level coordination in the absence of information about competition and others' presence could reflect benefits from diversity. This is in line with findings, that supposedly irrational individuals can produce group-rational behavior (Curşeu et al., 2013; Sasaki & Pratt, 2011), and uninformed and randomly acting individuals can benefit consensus and coordination (Couzin et al., 2011; Shirado & Christakis, 2017). Yet our results suggest that the benefits from such individuals can be especially pronounced in situations in which information about the social situation is limited. As soon as additional information is available, this may allow for superior strategies.

In sum, whereas, adaptation seems strongly characterized by individual reinforcement learning or robust strategies, information about others' behavior also appears to be integrated in the decision process. At least some individuals may adapt their behavior, if a situation is represented as social and competitive. Learning from individual reinforcement and social information may be complemented by the use of robust strategies or an adaptation of information use in response to assumptions about the situation.

### 2.9.3 Using computerized agents

Finally, our study highlights both the value and limitations of computerized learning agents (see also Schulze et al., 2015; Shirado & Christakis, 2017). The controlled setting provided valuable first insights into how individuals contribute to coordination in competitive environments. However, a fixed set of computerized agents does not allow inferences about interdependence between different individuals or about coordination patterns. Future research could further harvest the potential of computerized learning agents to disentangle the role of types of individuals for coordination. It is an interesting question what degree of stability on behalf of the agents is required so that groups manage to establish a stable IFD, even if this means selection of the less likely option.

### 2.9.4 Conclusion

Unsurprisingly, learning coordination under undisclosed competition is challenging. Nevertheless, individuals manage to establish some degree of adaptation when compared to non-competitive situations. The results account for the use of robust strategies but also suggest that different individuals may use different approaches to deal with this challenge. Future investigations should continue to disentangle the effects found here. Dependent on one's ability or willingness to engage with the setting, information about others' presence or behavior can improve adaptation to undisclosed competition. The present results point to various solutions that individuals have developed to deal with the challenge of learning interdependence under risk and uncertainty.

# 3

# Recommendation and exploration

## Providing and receiving recommendations affect exploration behavior in spatially correlated bandits

*Joint work with Hansjörg Neth.*[3,4] *and Wolfgang Gaissmaier,*[1,2,3,4]

[1] *Graduate School of Decision Sciences, Universität Konstanz*
[2] *Graduate School of the Social and Behavioural Sciences, Universität Konstanz*
[3] *Department of Psychology, Universität Konstanz*
[4] *Centre for the Advanced Study of Collective Behaviour, Universität Konstanz*

## Abstract

Relying on advice and recommendations provided by others can be helpful to navigate environments with large numbers of alternatives. Yet, the recommendation situation can affect information search both among individuals who provide recommendations (advisors) and individuals who receive recommendations (recipients). Here we use spatially correlated bandits, to investigate how advisors and recipients change their information search. Spatially correlated bandits were 11×11 grids, in which tiles are correlated so that individuals can generalize from their previous experience. First, we investigate advisors. They adjust their exploration behavior to find the best tile if they are rewarded according to recipients' performance. Second, we investigate how recipients respond to receiving a single tile as a recommendation. Previous findings imply a bias towards recommendations. We conducted a set of simulations, which indicate that in spatially correlated bandits this bias can reflect in (a) a large probability of taking unknown recommended options, (b) the adaptive dismissal of bad recommendations, (c) a preferential treatment of recommendations after taking them, resulting in a decreased probability of finding the best option. An experimental investigation shows that recipients' exploration behavior is strongly affected by receiving unknown recommendations, because they have a high likelihood of clicking on recommended tiles. Thereby, individuals generalize their previous experience to make the decision whether to dismiss a recommendation. The probability of finding the best option and distance to recommendations indicate that individuals largely do not treat recommendations preferentially after clicking on them and fail to exploit recommendations they receive. Together the results indicate that the recommendation situation can affect information search among advisors and recipients alike. We discuss implications for transmission chains of multiple advisors and recipients.

## 3.1 Introduction

We frequently rely on advice and recommendations to identify good options and solutions from an insurmountable number of alternatives. For instance, colleagues or supervisors provide us with helpful writing advice, or share their solutions to tasks like programming or data analysis. Social information based on others' experience is especially valuable if the environment is uncertain or features a large number of alternatives (Laland, 2004; Laland et al., 2020). Thereby, individuals often turn to experienced or prestigious individuals (Henrich & Gil-White, 2001; Jiménez & Mesoudi, 2019; Wood et al., 2013) or advisors (Littlepage et al., 1997) to learn about valuable alternatives.

Thus, many advice situations consist of an experienced advisor and a recipient who receives social information in form of advice or a recommendation from the advisor. Obviously, the advisor typically has explored at least some alternatives in the environment. Yet also the recipient may have previous experience, which allows them to judge the value of recommendations they receive.

Both advisors and recipients face the *exploration-exploitation dilemma* (Cohen et al., 2007; Hills et al., 2015; Mehlhorn et al., 2015). While the exploration of novel options carries the risk to miss out on certain gains from exploiting a known option, the exploitation of known options carries the risks to get stuck in an inferior option. For advisors this trade-off implies the decision, whether to search for a better recommendation or to increase their own certain gains. Recipients, in turn, face this dilemma twice: both if they explore the environment themselves and when they have to decide how much they want to rely on others' experience or their own (Frith & Frith, 2012). We ask, how the recommendation situation affects information search among advisors and recipients.

### 3.1.1 Literature review

How recipients generally use social information has been investigated within various research traditions. In studies of social learning (Kendal et al., 2018; Laland, 2004) and cultural transmission (Caldwell et al., 2020; Mesoudi, 2016), individuals can observe others, have access to others' solutions while they explore environments with multiple options (e.g., Efferson et al., 2008; McElreath et al., 2005), or design virtual artifacts (e.g., Acerbi et al., 2016; Caldwell & Millen, 2009; Derex & Boyd, 2016; Derex et al., 2013; Mesoudi, 2011). Conversely, studies on advice taking in judgment and decision making mostly relied on the Judge Advisor System (JAS, Sniezek and Buckley, 1995, see also Bonaccio and Dalal, 2006; Rader et al., 2017). A judge or advice recipient typically has the task to estimate or judge a quantity (e.g., the distance between two cities or the date of some event Yaniv, 2004). Then they receive the judgment of an advisor and can update their estimate.

Further findings indicate that the integration of social information may often be non-optimal and can have negative consequences. On the one hand, individuals rely heavily on social information and show strong conformity which can lead to the propagation of outdated information (Kameda & Nakanishi, 2002), maladaptive trends (Bikhchandani et al., 1998; Giraldeau et al., 2002) or herding behavior (Kameda & Hastie, 2015; Raafat et al., 2009). Similarly, in the JAS human individuals rely on advice even if they know that it stems from a less experienced advisor (Harvey & Fischer, 1997) or is of low quality or even random (Fiedler et al., 2019; Hütter & Fiedler, 2019; Schultze et al., 2017). On the other hand, individuals often rely too little on social information and weigh their own information more strongly and therefore forego potential benefits (e.g., Acerbi et al., 2016; Efferson et al., 2008; McElreath et al., 2005; Mesoudi, 2011; Mesoudi & O'Brien, 2008a). Likewise, when taking advice individuals show a consistent tendency for *egocentric discounting*, that is, they give more weight to their own information or opinion than to advice they receive (Bonaccio & Dalal, 2006; Rader et al., 2017; Yaniv & Kleinberger, 2000). Thus, individuals rely on social information including advice and recommendations but consider and integrate it—sometimes too much, sometimes too little.

These previous studies mostly addressed only either how individuals explore the available options in complex environments with observational information and teaching or to what degree individuals integrate advice in single estimates. As an exception, Biele and colleagues showed that individuals who received advice about an option in a multi-armed bandit task were more likely to click on that recommended option irrespective of its quality (Biele et al., 2009; Biele et al., 2011). They suggested that following advice

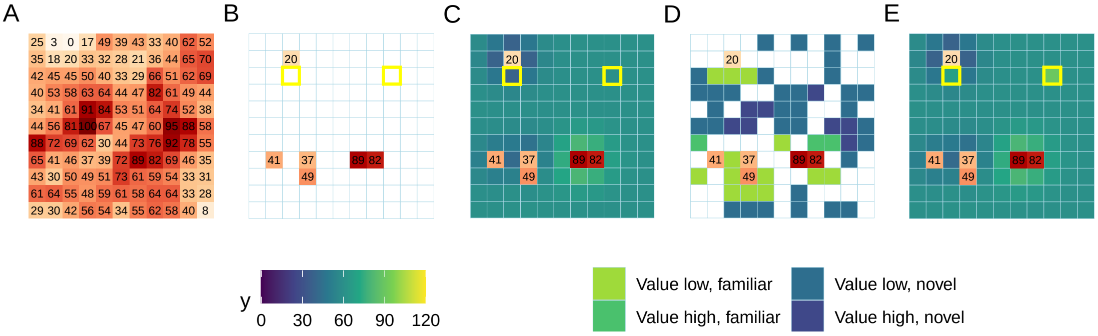

Figure 3.1: (A) Example of a spatially correlated bandit. (B) Possible advice taking situation (with 2 possible options highlighted in yellow) determined from (C) a utility landscape integrating estimated value and uncertainty with each other. (D) Illustration of options eligible for generating novel and familiar recommendations based on actual value and estimated familiarity (uncertainty). (E) Illustrates the estimated effect of a recommendation as a bonus to the estimated utility.

may be intrinsically rewarding and, thus, biases information search under uncertainty towards recommendations.

In sum, previous research indicates that social information and recommendations affect recipients' information search. Thereby, research on the search behavior of advisors has been scarce (Blunden & Gino, 2018; Bonaccio & Dalal, 2006; Jonas & Frey, 2003). Moreover, in many environments, generalization allows individuals to extrapolate their previous experience to unknown alternatives (see Gershman & Daw, 2017). Yet, it is unclear how recipients integrate recommendations if the environment allows them to generalize their previous experience. To address these gaps we investigate information search among advisors and recipients using spatially correlated bandits (Schulz et al., 2019; Wu et al., 2018).

### 3.1.2 Generalization and spatially correlated bandits

In spatially correlated bandits, individuals reveal tiles in a grid they cannot fully explore in the number of clicks they have (see Figure 3.1). Thereby, correlations between options allow individuals to inform their decisions by *function learning* (see panel A, e.g., DeLosh et al., 1997). That is, they can use previously encountered options to estimate a function over the values of all known and unknown options. In this task function learning has been shown to describe human behavior well (Schulz et al., 2019; Wu et al., 2018). This suggests that humans rely on generalization to inform their exploration behavior. Previous insights into individuals' behavior in non-social settings and the potential to assess effects of generalization make this task a good candidate to address our questions about effects of the recommendation situation on search behavior.

We rely on a combination of experiments and simulations in order to address our questions. First, we experimentally address how the prospect of providing a recommendation affects the search behavior of advisors, who have no direct incentive to deceive the recipients. This also allows us to gauge, whether advisors spontaneously pass on good recommendations so that relying on their information is a potentially valuable strategy for recipients. Second, we conduct simulations to assess, which strategies individuals can use to rely on recommendations and to derive predictions about potential behavior. Finally, we experimentally investigate how individuals integrate recommendations with varying levels of own previous experience.

## 3.2 Providing recommendations: Experiment 1

The prospect of providing a recommendation may affect exploration behavior and therefore the quality of recommendations. If advisors are motivated to give the best possible recommendation they should try to find the best possible option. Results from a non-social context indicate that the goal to find a large option affects exploration behavior Wu et al. (2018). In spatially correlated bandits, individuals whose goal it was to accumulate rewards explored more locally than individuals whose goal it was to find the largest possible option. Among those who were rewarded for finding the largest possible option less local exploration reflected in an increased Manhattan Distance (MD) to the previous option (that is the distance between tiles without diagonals). Similarly, it is to be expected that individuals who are rewarded according to the recipients' performance have an incentive to find the best possible recommendation. Exploration of a larger number of options increases their chances of finding the best option at no additional cost because they can pass on their best option anyways. Hence, advisors who are rewarded according to recipients' performance should explore more, behaving like individuals who are directly rewarded for finding the best option.

But what about individuals who are rewarded according to their own performance but face the prospect of providing a recommendation? They face more of an exploration-exploitation dilemma, because they may earn more from exploiting good options but this may lead them to end up with a worse recommendation. It is both conceivable that they try to actively explore more to increase their chances of finding a better recommendation and that they focus on their goal to accumulate rewards.

Some previous findings suggest that those who face the prospect of providing a recommendation explore more. Individuals who decide for others have been shown to search more information (Jonas & Frey, 2003). Moreover, Benjamin and Budescu (2015) showed that advisors sampled lotteries more frequently than individuals did in other sampling studies, outside of the advice context. One explanation they suggest is that the mere knowledge that others are affected leads individuals to searching more information. Unfortunately, they did not provide any direct no-advice comparison supporting their conjecture. We address this shortcoming and compare exploration behavior with and without providing a recommendation.

Since the concern for others' reward may affect information search, individuals' social preferences or social value orientation (e.g., Murphy et al., 2011) may affect their behavior. According to this measure, individuals can be classified according to how much they are concerned with others' rewards. Simply speaking, altruistic individuals sacrifice own rewards to increase others' rewards, prosocial individuals maximize the joint outcome, individualistic individuals maximize their own rewards, and competitive individuals actively decrease others' rewards. In fact, more altruistic individuals may be especially inclined to change information search, because they are most concerned with others' rewards. Similarly, prosocial individuals focus not only on their own benefit but also on others' benefit and the common good (Penner et al., 2005) and may, therefore, be more concerned by how their recommendation affects others.

Yet, the mere prospect of providing a recommendation may also be unrelated to a change in exploration. In spatially correlated bandits the strategies individuals use when trying to accumulate rewards may already balance exploration and exploitation (Wu et al., 2018). Thus, individuals will likely end up with good options even when they just explore as if they accumulate rewards.

In sum, individuals who are rewarded according to their own performance may search information differently when faced with the prospect to provide a recommendation but also might rely on the same strategies than when they don't. We want to shed light on under which conditions individuals are affected by the prospect of providing a recommendation.

### 3.2.1 Methods

We pre-registered our main predictions and the design under https://osf.io/cgvt8.

**Participants** We analyze the data from 154 participants (58 female, mean age 38.67 years ($SD$ = 12.16 years, range 19 to 70 years) recruited on Amazon Mechanical Turk (MTurk; https://www.mturk.com/). This sample size is lower than in our pre-registration, because our pre-registered criteria led the the exclusion of more participants than expected because of incomplete data or poor data quality, determined by self reports and repeated failure to correctly answer our understanding questions. Participants received a turn-up fee of US-$1.20 plus an average performance-based bonus of between US-$0.80 and 1.82 ($M = 1.54$, $SD = 0.17$). Deviating from the pre-registration, we did not perform additional participant exclusions, to retain acceptable statistical power.

**Materials** The experiment was conducted on unipark (http://www.unipark.com). The main task was implemented in custom JavaScript adapted from (Wu et al., 2018, https://github.com/charleywu/gridsearch/tree/master/experiment2D). The task environments consisted of 20 grids with $11 \times 11$ tiles each, which were selected from the environments used in (Wu et al., 2018, see Figure 3.1A). Environments were selected based on considerations related to Experiment 2.

**Design** In a within subjects design with 4 conditions we experimentally manipulated what individuals were paid for. In 2 conditions individuals were not asked to provide a recommendation and either received a performance-based reward for collecting as many points as possible (accumulation) or for finding the largest option (maximization). In the other 2 conditions individuals were asked to provide a recommendation and either were rewarded for collecting as many points as possible (accumulation plus recommendation) or according to the performance of other individuals who will receive the participant's recommendation (recipient performance). We manipulated between-subjects which condition individuals played first to assess potential order effects. Participants earned rewards by clicking on the tiles of the grid, each of which was associated with a reward value. Like in previous studies (e.g., Wu et al., 2018), the maximum value in each environment was drawn randomly to lie between 65 and 85 and other rewards were scaled accordingly. Thus, participants could not learn which maximum value to look for. Additionally, each option included a small amount of noise (an $SD$ of 1 point). The resulting payoffs were re-scaled to a range from 0 to 100 as such that participants could earn the same amount in each environment.

**Measures** Following Wu et al. (2018) the locality of exploration behavior indicating information search was measured as the average Manhattan distance (MD) of each click to the previously clicked option. The MD is defined as the distance between two tiles without going diagonally. To assess *prosocial orientation*, we used the six primary items of the Social Value Orientation (SVO) scale (Murphy et al., 2011). In each item, participants choose between nine different allocations for self and other. Based on the allocation choices over all nine items, participants can be categorized into one of four SVO categories: *Competitive*, *Individualistic*, *Prosocial*, and *Altruistic*. Further, we measured individuals willingness to take risks using the item "How do you see yourself: are you a person who is generally willing to take risks, or do you try to avoid taking risks?" (Falk et al., 2016; Falk & Hermle, 2018).

**Procedure** After providing informed consent, participants received task instructions and had to answer 3 understanding questions correctly before proceeding. In the main task each participant explored 8 grids, experiencing each of the 4 conditions twice. Prior to each grid they received instructions about their current goal and whether they would provide a recommendation. Moreover, before some randomly selected grids they were prompted to confirm their goal to ensure task understanding.

In those conditions, in which participants were asked to provide a recommendation they were shown their previous choices after finishing the grid. Then, they had to select an option to recommend to another person by clicking on the corresponding tile (they could also click on tiles they had not previously revealed). After the main task individuals answered the SVO scale and additional items, including the item measuring willingness to take risks. Eventually, individuals were debriefed about the study goals and informed about their payment.

**Data analysis** As pre-registered, we use linear multi-level regression models with random intercepts for individuals and environments to predict the average MD to the previous and recommended option in each round. We use the condition of the respective round as a predictor. In our main analysis, we contrast-coded the recommendation conditions. The first contrast compared the average MD of the two conditions in which

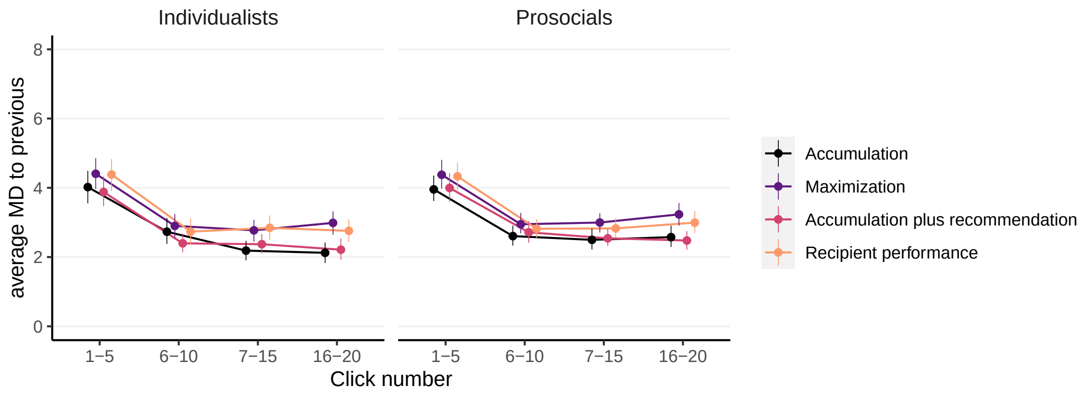

Figure 3.2: Locality of exploration as measured through the average MD to the previous option over the course of the round as a function of condition and social value orientation. Individuals who are rewarded for advice or are incentivized for finding the best possible option sustain less local exploration behavior. Error bars represent cluster-bootstrapped 95% confidence intervals.

participants provided a recommendation to the average MD when only accumulating rewards. The second contrast compared the two recommendation conditions against each other.

Beyond that, we explore effects of SVO type (effect coded, individualist vs. prosocial) and self-reported willingness to take risk ($z$-standardized) in a model including all 4 recommendation conditions (effect-coded). Additionally, we control for round number and condition order. The degrees of freedom for parameters are obtained with the Satterthwaite method in the package `lmerTest` (Kuznetsova et al., 2017). Post hoc tests were performed by comparing estimated marginal means with the package *emmeans* (Lenth, 2019), and degrees of freedom were obtained with the more conservative Kenward-Roger method.

### 3.2.2 Results

**Information search** Figure 3.2 illustrates that individuals explored less locally when they were rewarded according to advice recipients' performance. Thus, they behaved similar to individuals who searched for the local maximum. Those who provided a recommendation but were rewarded for accumulating rewards did not explore differently from those who were not asked to provide a recommendation.

In line with these observations, an effect of condition improves model fit, $\chi^2(2) = 28.906$, $p < 0.001$, $\Delta AIC = 24.91$. Neither including social value orientation, $\chi^2(1) = 0.223$, $p = 0.637$, $\Delta AIC = -1.78$, nor its interaction with condition into a model with all 4 conditions improves model fit, $\chi^2(17) = 0.502$, $p = 0.918$, $\Delta AIC = 5.50$ (for the model tables see Table B.2 models 1 and 2).

Planned contrasts excluding the maximization condition suggest that individuals explored less locally if they provided a recommendation, than if they were rewarded for accumulating points and did not provide a recommendation, $b = 0.065$, $SE = 0.023$, $t(751.71) = 2.837$, $p = 0.005$. The second contrast (difference between the recommendation only and the accumulation plus recommendation condition) suggests, that this effect is explained by the recommendation only condition, $b = 0.183$, $SE = 0.040$, $t(756.32) = 4.616$, $p < 0.001$ (for the model table see Table B.2 model 3). *Post-hoc* tests confirm that the difference between providing a recommendation and only accumulating rewards is driven by rounds in which participants were rewarded according to recipient performance, $\Delta b = -0.378$, $SE = 0.079$, $t(1061.65) = -4.793$, $p < 0.001$, but not by those who were rewarded according to their own accumulation performance, $\Delta b = -0.017$, $SE = 0.079$, $t(1059.63) = -0.221$, $p = 0.996$. Taken together, individuals do only seem to alter their exploration behavior if they are rewarded according to their recommendation, whereas the mere prospect of providing a recommendation does not lead to less local exploration.

**Passed on recommendations** Overall, in 74.84% of all rounds individuals passed on the best option they could see (85.21% passed on an option within 1 point of their best option). Surprisingly, in 12.18% of rounds in which a recommendation was passed on, individuals passed on options they did not select before. Excluding these individuals, still 85.21% passed on their best option, whereas 8.32% passed on an option worse than 5 points than their best revealed option. Thus, individuals did not always pass on their best option but mostly did so. Some individuals passed on options they did not know, which may be a symptom of careless responding, as well as of not finding any good option they found worth recommending.

**Recommendation quality** We do not find differences in the average value of recommendations. Whereas individuals in the accumulation plus recommendation selected options with an observed value of 68.61 ($SD = 6.61$), the value was 68.51 ($SD = 6.94$), when they were rewarded according to recipient performance, $t(153) = 0.151, p = 0.88$. Additionally, individuals did not differ in the value of their best option, irrespective of whether they provided a recommendation, 68.56 ($SD = 5.37$) or not, 68.94 ($SD = 6.46$), $t(153) = 0.016$, $p = 0.987$.

### 3.2.3 Discussion

We asked whether the prospect of providing a recommendation affects exploration behavior and therefore the quality of recommendations. We expected that advisors explore less locally if they know that they will provide a recommendation. Partially in line with our expectations, individuals only explored less locally if they were rewarded according to recipients' performance, but not if their goal was to accumulate rewards for themselves. Contrary to our expectations, this result was unaffected by whether individuals were more or less prosocial. Moreover, we find that most individuals passed on their best recommendation when asked to do so. Thus, it seems that advisors with potentially competing goals for exploration and exploitation resolve the exploration-exploitation dilemma by acting in their self-interest. Nevertheless, competing goals for information search do not appear to affect whether advisors pass on their best option.

Findings complement the scarce previous research on advice giving and information search (Blunden & Gino, 2018; Çelen et al., 2010). Individuals do not always seem to change their information search when making decisions for others (Jonas & Frey, 2003). If they perceive that information search could hurt their rewards, they appear to search less locally, like when they would not give a recommendation. Therefore, advisors' increased sampling of lotteries reported in Benjamin and Budescu (2015) is unlikely due to advisors knowing that they will affect others. In fact, their other explanations that the complexity of the lotteries or writing open ended descriptions required individuals to take more samples and that they enforced a minimum number of samples to be drawn appear to be more likely. However, their task and ours differ inasmuch that their advisors did not directly face competing goals except for a potential motivation to save time in sampling. What we can clearly state, is that advisors who do not face a potential conflict between exploration and exploitation change their search behavior, when they are asked to provide a recommendation.

Notably, in spatially correlated bandits it may not be maladaptive to resolve the exploration-exploitation dilemma rather egoistically, when accumulating rewards for oneself. Similar to Wu et al. (2018), we do not find that the more local exploration of accumulators harmed the quality of the best option and, hence, also does not harm advice quality. Given that the goal made little difference in the best option found, it may also be the case that accumulators are correct to not care about the consequences for advice recipients and focus on the task that pays off most for them.

The lack of losses in recommendation quality also explains why we do not find effects of our SVO measure. Whereas altruistic individuals would be expected to sacrifice own gains for others, prosocial individuals are expected to maximize *joint payoffs* (Murphy et al., 2011). Prosocial individuals were able to achieve this goal. In fact, exploitation maximized individuals payoffs and was not costly in terms of recommendation quality and exploration maximized both advisors' and recipients' rewards. Therefore, the absence of altruistic individuals in our sample may explain why we do not find effects of SVO. Thus, spatially correlated bandits are a setting in which exploration for others has little consequence and therefore many individuals behave according to their accumulation goal.

Nevertheless, in other settings which do require to decrease local exploration to find the best options, focus on an accumulation goal may lead to worse advice. Especially if joint maximization and maximization of own rewards conflict more strongly, we would also expect to find effects of social value orientation. Notably, if individuals then have a choice between an advisor who has searched options for their own benefit and one who has an interest in the recipients' performance, they should have a small preference for the latter. We encourage future research in different environments to investigate whether individuals adapt their behavior, or whether we then observe worse recommendations under conflict of interest.

## 3.3 Receiving recommendations

We now turn to investigate how receiving a recommendation affects recipients' exploration behavior. Therefore, we first identify factors that potentially influence how individuals take recommendations and how they shape their exploration behavior. Then we provide results from simulations that inform our predictions that we address with experimental results.

First, individuals' behavior should be informed by the fact, that they could potentially benefit most from novel information. Individuals who stay close to already revealed options may miss out on good options in unexplored areas. This is similar to estimation tasks, where uncorrelated information from multiple individuals is important that groups benefit from aggregation (Lorenz et al., 2011). Accordingly, individuals probably can benefit most from taking unknown recommendations that are far from what they know (*novel* recommendations).

Despite these potential benefits, individuals more likely show egocentric discounting and dismiss advice distant from their estimate or opinion in the JAS (Yaniv & Milyavsky, 2007). In a similar vein, when individuals make a judgment, they only weakly weight or ignore divergent opinions (Moussaïd et al., 2013) or beliefs (Giese et al., 2019; Moussaïd et al., 2015). Although in the case of recommendations this can be adaptive to eliminate extreme opinions (Harries et al., 2004), it can also lead individuals to miss out on beneficial novel information. Thus, although individuals can benefit from novel recommendations, it could be that they dismiss them more likely. Accordingly, they could be less likely to click on those recommendations and also be less likely explore close to them.

Another factor that could both negatively or positively affect individuals use of recommendations, is the recipient's experience. In the JAS it has been shown that more experienced individuals adjust their judgments less towards advice they receive (Yaniv, 2004). Similarly, Harvey and Fischer (1997) show that individuals more strongly rely on advice from experienced advisors. Yet, they also show that more experienced individuals are more likely to take advice in important decisions. In sum, more experienced individuals could be either more or less willing to take recommendations.

The environment allows to generalize previous experience. This could allow individuals to judge the value of a recommendation they have not revealed before. In line with this idea, it has been found that individuals are able to ignore bad advice (Gardner & Berry, 1995; Yaniv & Kleinberger, 2000) and to use individual learning to judge advisors expertise (Boorman et al., 2013) or intentions (Diaconescu et al., 2014) as indicators of advice quality. Thereby, more experienced individuals could be at an advantage, because they have more previous information to judge recommendations. However, there is also indication that individuals often fail to ignore bad advice (Hütter & Fiedler, 2019; Schultze et al., 2017), even despite superior knowledge (Fiedler et al., 2019). Therefore, we investigate whether individuals are able to use generalization to dismiss bad recommendations or whether they fail to do so.

Finally, it is an interesting question, whether recommendations are treated differently from self-selected options. As suggested previously, following recommendations can be intrinsically rewarding (Biele et al., 2009; Biele et al., 2011). Accordingly, additional utility from clicking on recommendations could affect subsequent information search by creating a bias towards the recommended option. In fact, such a bias together with

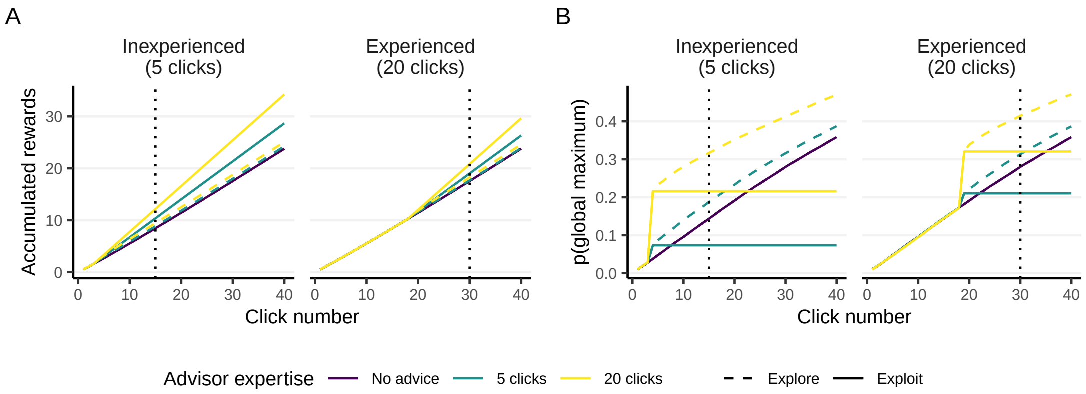

Figure 3.3: Simulated performance when taking, and exploring or exploiting recommendations from advisors with different expertise (columns). (A) The average cumulative reward is larger than without advice or continued exploration, whether individuals follow advisors with 5 or 20 clicks of experience. Conversely, (B) the proportion having found the global optimum is largest if individuals continue exploration and is even lower than without advice, dependent on how long they could continue exploration. Simulations are based on 40,000 individuals in each shown conditions (they can further be distinguished by 40 environments and differ in their parameter values).

failure to ignore bad recommendations could lead individuals to rely more strongly on bad recommendations than they should. In a related vein, children have been shown to explore the functions of an object less if they have received instruction on it (Bonawitz et al., 2011). Consequently, individuals could become more likely to exploit a recommendation or explore adjacent alternatives than to continue exploration.

On the contrary, in experiments with multi-modal landscapes (with multiple locally optimal peaks) it has been shown that individual learners without social information are more likely to get stuck in locally optimal options than social learners (peaks lower than the highest peak, e.g., the right peak in Figure 3.1A, Mesoudi and O'Brien, 2008a, 2008b). This difference was shown to be most pronounced, if the best options are narrowly defined, so that hill-climbing strategies are difficult (Acerbi et al., 2016; Acerbi et al., 2011). Yet, a bias towards recommendations could increase the chances that individuals rather resolve the exploration-exploitation trade-off towards exploitation of a recommendation. In this case, it could be social learners who become more likely to get stuck with a recommendation, putting them at risk to miss out on superior options. To shed additional light on the potential role of the above factors in spatially correlated bandits, we conducted a set of simulations.

### 3.3.1 Simulation results

Simulations that compare ignoring (continuing exploration) and following (exploiting) recommendations show that if individuals' goal is to accumulate rewards just blindly taking and following can be a simple, yet robust, heuristic strategy. We simulated individuals who explored the environment according to a function learning model and environments from Wu et al. (2018). Simulated individuals explored the spatially correlated bandits and received a recommendation after 5 or 20 trials ("clicks"). Then, they

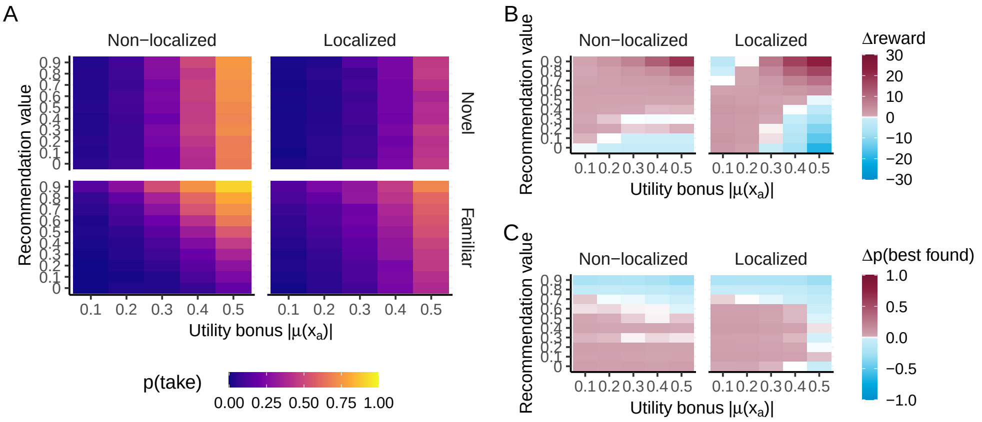

Figure 3.4: Results from a function learning model with utility bonus. (A) The probability of selecting a recommendation increases with the value of the utility bonus. For familiar options it additionally increases with recommendation value, suggesting that the function learning model generalizes from previous experience to make selecting good recommendations more likely. (B) Relative to receiving no recommendation, the average rewards after receiving a recommendation Δreward are increased for high utility bonus and high recommendation value but decreased for low recommendation value. Conversely, (C) the probability of finding the best option $\Delta p(\text{best found})$ decreases with utility bonus and high option value, whereas it is increased for low option value.

either continued exploration or exploited the recommendation. The simulation findings are summarized in Figure 3.3 (further details can be found in Appendix B.1).

If individuals' goal is to accumulate rewards, they can benefit from exploiting the recommendation. In fact, their average rewards are never lower than without a recommendation (colored lines, Panel A). However, exploitation comes at the cost of a decreased chance of finding the best option (Panel B). Accordingly, a simple heuristic that always takes and exploits a recommendation can already be expected to be fairly robust with respect to accumulating rewards and also finding the global optimum for short time horizons. Given that in uncertain environments simple heuristic strategies often perform well (Gigerenzer et al., 1999; Gigerenzer & Gaissmaier, 2011), individuals could rely on a simple "take-unknown-recommendations-and-exploit"-heuristic.

Instead of such very simple heuristics, individuals may also generalize from previous experience. This could play a role both, when they decide to take a recommendation or to integrate the recommendation afterwards. We rely on previous research by Biele et al. (2009), which suggests that clicking on recommendations may be intrinsically rewarding. This intrinsic reward can be thought to reflect a utility bonus to a recommended option (see also Figure 3.1). The combination of such a utility bonus with a function learning account yields a model with a bias for recommended options but still integrates them with previous experience (for details on the model see Appendix B.1).

Crucially, exploration according to such a model may additionally be subject to localization (Wu et al., 2018). This means that individuals prefer to stick close to their previous choice. This is achieved by decreasing the probability of selecting an option with its distance from the previous choice. Since this may affect how recommendations are used, we consider both localized and non-localized variants of the model.

We conducted a second set of simulations with this model to address how such a bonus can affect taking and exploration behavior. Again, simulated individuals acquired previous experience of 5 or 20 clicks. Then, starting from this situation, they received one randomly determined, unrevealed tile as recommendation. Recommendations were unrevealed tiles that were either familiar or novel, that is, close to known options or not.

Figure 3.4A shows, that the probability of selecting a novel option only increases with utility bonus ($|\mu(\mathbf{x_a})|$). The probability of clicking familiar options increases additionally with recommendation value. Thus, a function learning model with a bias towards recommended options can adaptively identify good familiar recommendations before taking them. Conversely, the model predicts an equal probability for selecting high and low novel recommendations.

Figure 3.4B and C show, how performance is affected for the same individuals across the 10 clicks after the recommendation is received. Like our previous simulations indicate, individuals with a large option bonus (high chance of clicking on a recommendation) accumulate more rewards than individuals without recommendations (Panel B). However, the largest benefits for average rewards coincide with decreased chances of finding the best option (Panel C).

In sum, we conjecture that individuals could rely on simple heuristic strategies but also on strategies that are biased towards clicking on recommendations. The latter allow them to adaptively dismiss some bad options outright. Interestingly, the same parameters that produce this adaptive taking behavior also lead to a large probability of sticking close to that option, hence, increasing the chances that the global optimum is not found. It is an open question, how individuals actually handle this situation. Therefore, we next experimentally investigate predictions informed by heuristic and function learning models.

### 3.3.2 Experiment 2

How do recommendations actually affect recipients' exploration behavior? Our simulations suggest that recipients can benefit from taking recommendations from experienced advisors. If recipients' goal is to accumulate rewards, they benefit most from exploiting the recommendation. At the same time, this decreases their chances of finding the best option. To provide insights into individuals' behavior, we experimentally investigated, (a) whether they indiscriminately take recommendations they receive or use their previous experience to select good over bad recommendations, as well as (b) how they continue exploration after receiving a recommendation and whether they exploit good and adaptively ignore bad recommendations.

Crucially, the experimental results on the generation of recommendations indicate that the assumption in the simulations that individuals typically receive good recommendations is justified. Advisors do not explore differently, when providing a recommendation after accumulating rewards for themselves. Since our simulations precisely assumed this, the advisor-side assumptions that underlie our simulations seem to be largely met. Hence, we inform our expectations for how human recipients deal with recommendations by means of our simulations and previous research.

**Taking recommendations** Although it may seem obvious, the decision to take a recommendation is the first possible influence on exploration behavior. Instead of exploring based only on previous experience, individuals may be pulled towards the recommendation.

If recommendations are provided by experienced advisors, always taking recommendations may be a good strategy. However, in real-world settings individuals may also have to deal with bad recommendations. In Experiment 1 we observe that individuals sometimes even pass on bad recommendations, although they have no incentive to do so. Generalization from previous experience allows recipients to judge recommendations before clicking them. This can help them to avoid wasting time with bad advice. If individuals rely on strategies that can be described by biased function learning, we should observe a high probability of clicking on recommendations and a higher probability of clicking on high, rather than low familiar recommendations. Conversely, if individuals deterministically click on or ignore recommendations, we should not observe the latter effect.

**Integrating recommendations** A bias for recommendations could also have a more sustained influence on recipients' exploration behavior. In fact, it could pay off to treat recommendations differently, after they are received. Simulations suggest, that individuals who want to accumulate rewards, can benefit from exploiting recommendations. Yet, this increases the likelihood that they get stuck in non-optimal options if they receive a recommendation that is not the best option (global optimum). To asses the impact of a bias towards recommendations, we investigate individuals' likelihood of finding the best option.

The distance to the recommendation serves as an additional measure of its impact beyond the likelihood of finding the best option. In the JAS the impact of advice is typically assessed by measuring, how much individuals adjust their initial estimate after receiving advice (see Bonaccio & Dalal, 2006). In a similar way, spatially correlated bandits allow to use the distance to the recommended tile as a measure for how much the recipient's exploration behavior is affected by the recommendation. Thus, the closer a recipient explores to an option, the more they can be thought to be affected by it.

Moreover, we investigate two additional influences on individuals' use of recommendations: recipient's experience and predictability of the environment. Whereas our simulations did not show differences due to recipient experience previous research suggests that the more experienced may be less likely to take recommendations and explore the effects on exploration behavior due to egocentric discounting (Harvey & Fischer, 1997; Yaniv, 2004). Therefore, we address recipient experience nevertheless.

Finally, also the structure of the environment could matter for behavior and outcomes. Wu et al. (2018) used rough and smooth spatially correlated bandits, which are defined through the degree of correlations between options. Whereas smooth environments have well defined peaks, rough environments feature more peaks and bad options have a somewhat higher likelihood to be close to good options. On the one hand, exploration close to recommendations can make individuals more likely to get stuck in some peak in smooth environments. On the other hand, previous research suggests an advantage of social learning in similar environments (Acerbi et al., 2016; Acerbi et al., 2011; Mesoudi & O'Brien, 2008b). To account for these potential influences we additionally investigate the role of environment type.

In sum, we use predictions based on a function learning model with a bias for recommendations to investigate how recommendations affect recipients' exploration behavior. Specifically, we first address, how a bias towards recommendations affects individuals' exploration behavior through increasing the probability to click on recommended tiles. We expect to find (a) a bias towards recommendations, reflecting in an increased probability of clicking on recommended tiles, and (b) an increased probability of clicking on recommended tiles for high-, as compared to low-valued, familiar recommendations, which indicates, that individuals adaptively use previous information to guide their taking decision. Secondly, we ask whether a bias for recommendations has sustained effects on exploration behavior beyond clicking on recommended tiles. We expect to find that (c) a decreased probability of finding the best option reflects that individuals treat recommended tiles differently from self-selected options and get stuck in recommended options due to their intrinsically rewarding nature, as well as (d) a bias towards recommendations reflects in a decreased distance to recommended options.

### 3.3.3 Methods

The predictions and design, as well as analyses were preregistered under https://osf.io/msre5. Deviations from this pre-registration are pointed out in the text.

**Participants** We analyze the data from 353 participants (121 female, mean age 37.12 years, $SD = 10.17$ years, 18 to 67 years) recruited on Amazon Mechanical Turk (MTurk; https://www.mturk.com/). This sample size is lower than the 400 participants aimed at in our pre-registration, because we had to exclude more participants than expected because of incomplete data or poor data quality. Exclusions were pre-registered and determined based on self reports and repeated failure to correctly answer our understanding questions, as detailed in our preregistration. We technically ensured that participants who took part in the advice giving study could not take part in the advice taking study. Participants received a turn-up fee of US-\$1.00 plus an average performance-based bonus between US-\$1.31 and 2.56 ($M = 1.81$, $SD = 0.21$). Deviating from the pre-registration, we did not perform additional participant exclusions, to retain acceptable statistical power.

**Materials** Like in the study on providing a recommendation the experiment was conducted on unipark (http://www.unipark.com) and the main task was implemented in custom JavaScript adapted from Wu et al. (2018). Each individual played 12 different task environments (grids). The 12 task environments were randomly drawn from 20 grids with $11 \times 11$ tiles each, selected from the grids used by Wu and colleagues. Environments were selected based on yielding the largest effects on finding the global maximum in a set of simulations and the constraint that we wanted 10 environments with high spatial correlations (smooth) and 10 environments with lower spatial correlations (rough). They were the same environments that had been used in Experiment 1, so that we could provide actual recommendations.

**Design** We experimentally varied the recommendation quality (high vs. low) and familiarity (familiar vs. novel), as well as recipient experience (5 vs. 15 clicks prior to the recommendation) in a $2 \times 2 \times 2$ within-subjects design. We selected *unrevealed* recommendations unknown to the participants from (a) one of 2 ranges of option quality (low: 40 to 60 points and high: 75 to 99) and (b) one of 2 ranges of uncertainty (familiar: 0 to 0.9 and novel: 0.9 to 1, estimated according to a function learning algorithm). These generated recommendations never included the global optimum. This maximizes the chances that we observe effects on finding the best option if individuals get stuck in high-valued recommendations. Additionally, individuals received 2 control rounds without any recommendation and 2 rounds with genuine recommendations provided by other individuals. Finally, we randomly varied the environments individuals explored because we did not have clear hypotheses about their effects.

To avoid deception, we used a conditional information lottery design, which has previously been used in the context of artificially generated social information (Bardsley, 2000; Morgan et al., 2012). Participants were informed that there were only 2 rounds with genuine recommendations, which are relevant for their payment but were not informed, which rounds they are. It has been shown that individuals treat all information as genuine under these conditions (Bardsley, 2000). Additionally, participants were informed that they were also paid for 1 of 2 rounds without recommendation, to ensure that their effort would not differ between rounds with and without recommendations because rounds without a recommendation would be paid with certainty.

Individuals played for a total 15 or 25 clicks in each round. They earned rewards by clicking on the tiles of the grid, each of which was associated with a reward value. Like in previous studies (e.g., Wu et al., 2018), the maximum value in each environment was drawn randomly to lie between 65 and 85 and other rewards were scaled accordingly. Thus, participants could not learn which maximum value to look for. The resulting payoffs were scaled back to a range from 0 to 100 as such that participants could earn the same amount of US-\$0.65 in each of the 3 rewarded grids. Additionally, each option included a small amount of noise (an $SD$ of 1 point).

**Measures** To assess effects of *recommendation value* we use the respective option's value relative to the environment maximum. Consequently, 1 corresponds to the globally best option and 0.5 indicates an option that is half the size of the globally best option. In order to distinguish effects of mere option value and effects of treating advice differently, we use a control option in rounds without a recommendation. To this end, we use the option clicked after a recommendation *would have been received* as control option. That is the option 9 clicks before the end of the round (click number 6 for 15 clicks and click number 16 for 25 clicks). We will refer to the value of this control option as recommendation value as well.

The locality of exploration behavior was measured as the average MD of each click after the recommendation to one of two reference points: (1) The recommended option or, in case no advice was received, the option clicked in the trial of the recommendation and (2) the previously clicked option. More exploratory behavior can be expected to reflect in less local exploration and, hence, in larger MDs in both cases. We also included additional measures that are not discussed in the present paper.

**Procedure** After providing informed consent, participants learned about their task to explore environments and accumulate as many rewards as possible. They were informed that tiles are spatially correlated and that they can infer the quality of neighboring tiles. Before they started the main task, they were asked 3 questions to ensure their understanding of the task, they had to answer correctly to continue. After that, each participant explored a total of 12 different environments, 1 under each of $6 \times 2$ conditions (4 times generated advice, 1 time genuine advice, 1 time no advice, each for 15 and 25 clicks respectively). After that they answered some open questions about how they used the recommendations and our additional measures. Finally, participants provided demographic information and were debriefed about the study goals.

**Data analysis** As preregistered, we model our data on the level of rounds as unit of observation using multi-level models estimated with `lme4` (Bates et al., 2015) in R (R Core Team, 2018). To control for non-independence between individuals and cross-classification of environments, we test random intercepts for individuals and environments (Baayen et al., 2008; Gelman & Hill, 2007). In order to “keep it maximal” (Barr et al., 2013), we also test random slopes for our within-subjects predictors.

We investigate our dependent variables on the level of the rounds that participants played. To investigate advice taking, we investigate the probability of clicking the recommendation (immediately after receiving it and over the course of the round). To investigate advice integration, we predict the probability of finding the best option, as well as the average MD to the the recommendation.

We use logistic mixed regression models to predict the probability of clicking the recommendation and the probability that the best option is found. To predict the average MD to the recommended option we use linear multi-level regression models. As predictors we test the value and familiarity of the recommendation, as well as recipient experience and environment type. In all analyses, except for the probabilities of taking, we additionally include two dummy coded indicators for (a) whether no recommendation or (b) a genuine recommendation was received. Here we deviate from our pre-registration, because this allows us to better address and understand differences of generated recommendations to self-selected options and genuine recommendations. We enter predictors and their interactions block-wise and report the fits of a best fitting model. Model tables and details on the selection procedure can be found in Appendix B.3.

For comparing models we report likelihood ratio tests and differences in AIC ($\Delta AIC$), where positive differences indicate evidence in favor of a model. Here, we deviate from our pre-registration and test all fixed effects and their interactions, and then test their random slopes, to arrive at more parsimonious models. For testing fixed effects we rely on likelihood ratio tests performed by the function `Anova` in the package `car` (Fox & Weisberg, 2019), testing parameters by dropping each in turn. To further evaluate differences between conditions we use post-hoc tests of the estimated marginal means conducted with the package `emmeans` (Lenth, 2019).

The degrees of freedom for parameters in linear models are obtained with the Satterthwaite method in the package `lmerTest` (Kuznetsova et al., 2017). Across analyses the recommendation value and the current best option are centered to a value of 0.65 and scaled to steps of 0.15. The effects of recommendation familiarity (novel vs. familiar), recipient experience (5 vs. 15 clicks before the recommendation), and environments (smooth vs. rough) are included as effect-coded variables.

### 3.3.4 Results and discussion

#### 3.3.4.1 How is exploration behavior affected by clicking on recommended tiles?

First, we consider the probability that individuals take a recommendation. Across our 4 conditions with generated recommendations the recommended option was clicked in 86.44% of all rounds. But only in 67.46% of all rounds the recommendation was clicked immediately after it was received. Thereby, 54.96% of all individuals clicked on the recommendation in all 8 rounds. However, only 22.95% did so directly. Of all individuals 21.81% never clicked on any recommendation. Although the overall propensity to click on a recommendation is high, few individuals clicked on all recommendations. This accounts against a strict “always take advice”-heuristic among a large proportion of individuals.

Thus, receiving a recommendation affects exploration. In line with a bias towards recommended options, individuals are likely to click unknown recommended tiles. But do individuals adaptively use previous information to inform their decision? To this end we consider first the probability that individuals click on a recommendation immediately after receiving it and then whether they click it at all over the course of a round.

**Which factors affect clicking on recommendations immediately?** First, we consider the results of a generalized logistic mixed model predicting the probability that individuals click on a recommendation immediately after receiving it (for details on the model selection procedure see Appendix B.3). The results of the best fitting model are displayed in Figure 3.5A. The following effects ($b$-values) are on the logit-scale.

Contrary to our expectations, we do not find evidence that individuals rely on their previous experience to decide whether to take or dismiss a recommendation immediately. We do not find evidence for an interaction of recommendation value and familiarity, $b = 0.017$, $SE = 0.014$, $OR = 1.017$, Wald $Z = 1.155$, $p = 0.248$, or its higher order interactions. Yet, a main effect of recommendation familiarity provides evidence that individuals click more likely on familiar recommendations, $b = 0.157$, $SE = 0.049$, $OR = 1.169$, Wald $Z = 3.197$, $p = 0.001$. This familiarity effect is decreased in rough environments and among the more experienced, as indicated by 2-way interactions of familiarity and environment type, $b = -0.106$, $SE = 0.051$, $OR = 0.899$, Wald $Z =$

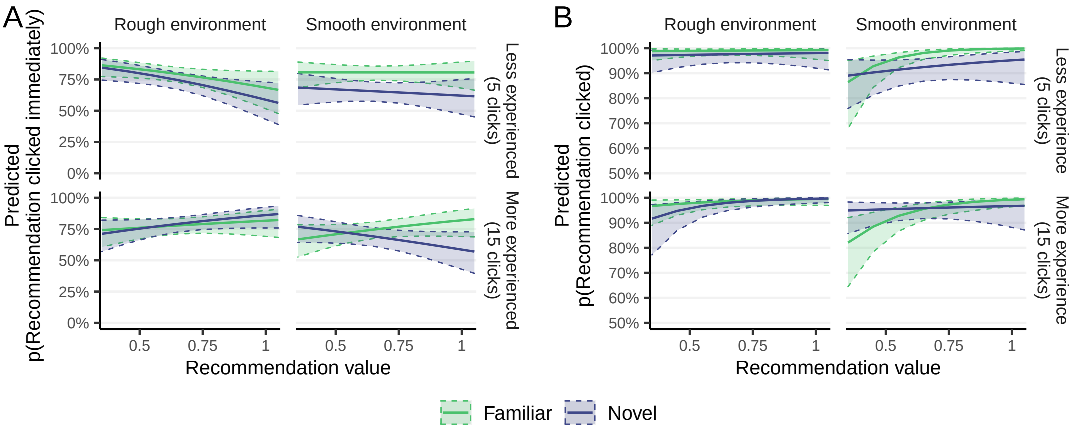

Figure 3.5: Estimated probability of clicking on the recommendation as a function of recommendation value, familiarity, environment type (rough vs. smooth) and recipient experience (5 vs. 15 clicks before receiving a recommendation) (A) immediately after receiving it and (B) anytime within a round. Individuals generally prefer familiar recommendations. Results on immediate clicking are mixed. Yet, when considering clicking anytime in smooth environments individuals are more likely to dismiss low valued familiar recommendations.

$-2.069$, $p = 0.039$, and $b = -0.108$, $SE = 0.049$, $OR = 0.898$, Wald $Z = -2.198$, $p = 0.028$. Together with a positive main effect for rough environments this indicates that individuals rely more on both types of recommendations if environments are less predictable.

A positive 3-way interaction of recommendation value, recipient experience and environment type, $b = 0.034$, $SE = 0.015$, $OR = 1.034$, Wald $Z = 2.211$, $p = 0.027$, additionally indicates that individuals vary with regard to which recommendations they click on. Surprisingly, less experienced individuals in rough environments and more experienced individuals in smooth environments are less likely to click on high valued novel recommendations (see Figure 3.5A). That is, they appear to systematically miss out on higher valued novel recommendations. How this phenomenon can be explained is unclear. We can largely exclude that it is due to more negative expectations for these recommendations, because in both cases the correlation between recommendation value and estimated value from a function learning model are positive. Potentially, it is due to an interplay of preferring close recommendations and not acting on current estimates. At least it could indicate the maladaptive nature of a bias for familiar recommendations.

In sum, we don't find evidence that individuals (at least not successfully) use previous information to determine whether to immediately click on high-value recommendations more likely. Yet, they prefer to click on recommendations close to revealed tiles. Thereby, they do not make this difference in less predictable environments.

**Which factors affect clicking on recommendations over the course of a round?** Whereas we do not find evidence that individuals use previous information about option values to inform their decision to click immediately on a recommendation, we find indication that individuals may inform their decision to ignore recommendations in more predictable environments. Again, we consider the results of a generalized logistic mixed model predicting the probability that an individual clicks on the recommendation anytime during a round (for details on the model selection procedure see Appendix B.3). Figure 3.5B summarizes the model results. Again, the following effects ($b$-values) are on the logit-scale.

A main effect of recommendation value indicates that individuals on average are more likely to click on higher-valued recommendations, $b = 0.146$, $SE = 0.054$, $OR = 1.157$, Wald $Z = 2.671$, $p = 0.008$. Moreover, a marginal main effect of familiar recommendations, $b = 0.333$, $SE = 0.173$, $OR = 1.395$, Wald $Z = 1.927$, $p = 0.054$, indicates that individuals on average are more likely to click on familiar recommendations. We were interested, if individuals rely on previous experience to decide which recommendations to click. The expected 2-way interaction of recommendation value and familiarity is only marginal but in the expected positive direction, $b = 0.050$, $SE = 0.027$, $OR = 1.051$, Wald $Z = 1.843$, $p = 0.065$. Crucially, the 3-way interaction of recommendation value, familiarity and environment indicates that the difference in the increase with recommendation value between novel and familiar recommendations additionally varies by environment type, $b = -0.076$, $SE = 0.025$, $OR = 0.927$, Wald $Z = -2.986$, $p = 0.003$. The difference in increases is estimated to be smaller in rough environments and hence is larger in smooth environments. Comparisons of simple slopes by environment and recipient experience indicate that the positive effect of recommendation value on cumulative taking is larger among familiar recommendations as compared to novel recommendations, both when individuals were less experienced, $\Delta b = 0.282$, $SE = 0.102$, $OR = 1.325$, Wald $Z = 2.749$, $p = 0.006$, and when they were more experienced, $\Delta b = 0.223$, $SE = 0.096$, $OR = 1.249$, Wald $Z = 2.328$, $p = 0.020$. We take this as evidence that individuals use previous information to decide whether to ignore a recommendation if they explore more predictable smooth environments.

Interestingly, an interaction of familiarity with recipient experience shows that the preference to click on familiar recommendations is less pronounced among the more experienced, $b = -0.261$, $SE = 0.079$, $OR = 0.770$, Wald $Z = -3.322$, $p = 0.001$. A main effect of recipient experience is not systematically different from zero, $b = -0.056$, $SE = 0.077$, $OR = 0.946$, Wald $Z = -0.726$, $p = 0.468$. Thus, contrary to expectations, the more experienced are not less likely to click on recommendations overall. Moreover, they show a decreased familiarity effect relative to the less experienced, indicating that they use novel information more than less experienced individuals.

In sum, the decision whether individuals click on a recommendation anytime during a round is affected by its familiarity. Individuals are especially more likely to immediately click on familiar recommendations. Nevertheless, the probability that they click on a recommendation during a given round is very large, indicating that they largely change their locality of exploration in favor of taking the recommendation. Yet, the recommendations that remain unclicked are typically either familiar and of low value or novel, especially if individuals are less experienced. Thus, individuals appear to rely on a combination of strategies that (a) mostly result in clicking on a recommendation, (b) discount novel information and (c) incorporate previous experience to decide which recommendations to take, at least in more predictable environments.

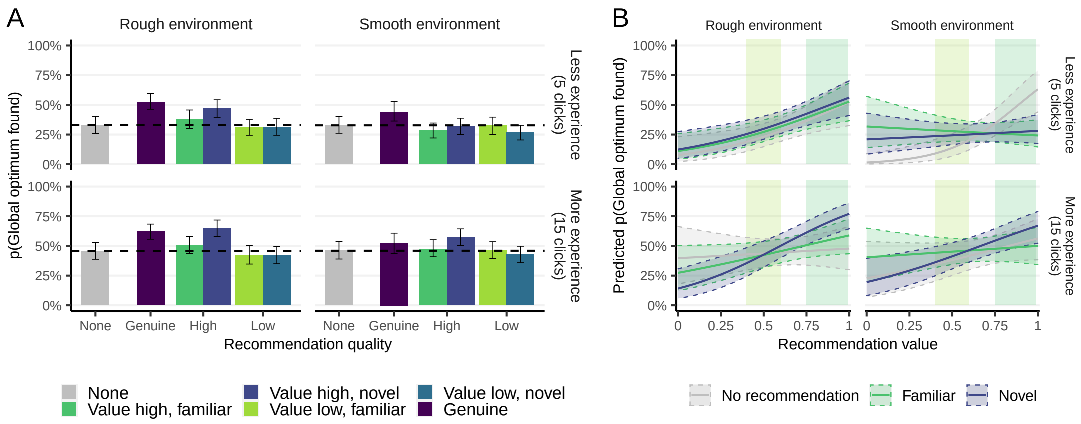

Figure 3.6: The probability of finding the best option (A) as round-level proportions in the data and (B) estimated from a generalized logistic mixed model as a function of condition (value high vs. low and familiarity familiar vs. novel), environment type and recipient experience. The (A) dashed and (B) grey line indicate the probability of finding the best option if no recommendation is received. (A and B) High-valued novel options provide an advantage to finding the global optimum relative to receiving no recommendation or a low value recommendation. (B) Less experienced individuals in smooth environments are less likely to find the global optimum if they receive a recommendation than if they receive a recommendation of similar value. Error bars and confidence bands correspond to 95% confidence intervals. In (A) confidence intervals were cluster-bootstrapped on the individual level. Shaded regions in (B) indicate the lower and upper ranges of generated recommendation values respectively.

#### 3.3.4.2 How do recommendations affect further exploration?

Next, we address how recommendations further affect exploration behavior. If individuals restrict their exploration behavior, their probability of finding the global optimum and the (Manhattan) distance (MD) to the recommendation should be decreased. We do not find consistent evidence for such a restriction of exploration across conditions but in limited situations with less experienced recipients in smooth environments.

**How is the probability of finding the best option affected?** To answer whether individuals risk getting stuck, when they receive good recommendations, we investigate how recommendations affect individuals' likelihood of finding the best option. Contrary to our expectations, Figure 3.6A shows that high-valued and genuine recommendations rather *increase* the probability that individuals find the global maximum relative to receiving no recommendation. Only among rounds with low-valued recommendations and among less experienced individuals who receive high-valued recommendations in smooth environments the chances of finding the global maximum tend to be lower than if they receive no recommendation (dashed line). This could be an indication that less experienced individuals get stuck in bad recommendations.

We further corroborate our results with the results from a generalized logistic mixed model predicting the probability of finding the globally optimal option (for details on the model selection procedure see Appendix B.3). This model allows to compare recommendations to corresponding self-selected options by including the value of the option selected after a recommendation would have been received in no-recommendation rounds. Since the no-advice comparison (and the genuine comparison) are dummy coded, effects without them pertain to generated recommendations only. Figure 3.6B summarizes the model results on the probability scale. Again, the following effects ($b$-values) are on the logit-scale.

Are there conditions under which individuals actually get stuck? A positive effect of recommendation value (and its no-recommendation control), $b = 0.071$, $SE = 0.013$, $OR = 1.073$, Wald $Z = 5.629$, $p < 0.001$, clearly indicates that the best option is found more likely, the higher the value of the recommendation received. A marginal negative difference to no-recommendation rounds, $b = -0.179$, $SE = 0.102$, $OR = 0.836$, Wald $Z = -1.754$, $p = 0.079$, indicates decreased chances to find the global optimum if no recommendation is received. Absence of evidence for an interaction of receiving no recommendation with the control value, $b = 0.044$, $SE = 0.028$, $OR = 1.045$, Wald $Z = 1.566$, $p = 0.117$, shows that this difference does not vary with recommendation value. An increased slope of recommendation value for genuine recommendations (not shown in figure) reflects high-valued genuine recommendations increase the probability of finding the global optimum, $b = 0.116$, $SE = 0.025$, $OR = 1.123$, Wald $Z = 4.702$, $p < 0.001$. Crucially, among genuine recommendations 31.87% were the global optimum. Nevertheless, these results suggest against a higher likelihood of getting stuck when high-valued recommendations are received.

The effect of recommendation value varies with recommendation familiarity among generated recommendations, $b = -0.026$, $SE = 0.013$, $OR = 0.974$, Wald $Z = -2.101$, $p = 0.036$, and environment type, $b = 0.039$, $SE = 0.013$, $OR = 1.040$, Wald $Z = 3.040$, $p = 0.002$. Effects indicate that the advantage of higher recommendation value is decreased among familiar recommendations and increased in rough environments. Moreover, the effect is modulated by 3-way interactions of recommendation (or no recommendation control) value with no recommendation rounds and higher recipient experience, $b = -0.084$, $SE = 0.028$, $OR = 0.919$, Wald $Z = -3.043$, $p = 0.002$, and no recommendation rounds and less predictable environments, $b = 0.766$, $SE = 0.265$, $OR = 2.150$, Wald $Z = 2.894$, $p = 0.004$. These interactions indicate a lower increase with the value of no recommendation control options among the more experienced and in rough environments. Put differently, among less experienced individuals in smooth environments the probability of finding the global optimum increases more strongly with the value of self-selected options. This, conversely, could indicate a disadvantage if individuals receive a high-valued recommendation under these conditions (see Figure 3.6B).

To understand the nature of these interactions better, we conducted additional comparisons. Specifically, we investigate the differences between a familiar or novel recommendation and receiving no recommendation for low (0.5 of the optimal option) options and high (0.8 of the optimal option) options respectively. Negative differences among the less experienced in smooth environments for novel, $b = -0.606$, $SE = 0.250$, $OR = 0.545$, Wald $Z = -2.429$, $p = 0.015$ and familiar recommendations, $b = -0.655$, $SE = 0.252$, $OR = 0.520$, Wald $Z = -2.594$, $p = 0.009$, indicate that self-selecting a high-valued option is related to higher chances of finding the global optimum than receiving it as a recommendation. Moreover, we find positive differences for low-valued familiar recommendations among the less experienced in smooth environments, $b = 0.886$, $SE = 0.401$, $OR = 2.424$, Wald $Z = 2.209$, $p = 0.027$, and high novel recommendations among the experienced in rough environments, $b = 0.886$, $SE = 0.401$, $OR = 2.424$, Wald $Z = 2.209$, $p = 0.027$. Differences suggest that individuals benefit from recommendations as compared to self-selected options when the probability of finding the maximum has been low or when the recommendation provides novel information. Thus, finding the best option is sometimes hindered and sometimes benefits from receiving recommendations.

Overall, individuals are not generally less likely to find the best option if they receive high-valued recommendations. They benefit most from good options in less predictable rough environments, especially, if recommendations are novel. Benefits of receiving a recommendation for finding the best option are smaller in smooth environments. In fact, less experienced individuals face potential disadvantages relative to receiving no recommendation. This could indicate that these individuals are more likely to get stuck in inferior options.

We also explored, whether local optima increase individuals' risk of getting stuck. Inclusion of an indicator whether an option was a local optimum improves model fit, $\chi^2(1) = 131.086$, $p < 0.001$, $\Delta AIC = 129.09$. Yet, its interaction with the difference between recommendations and self-selected options does not, $\chi^2(1) = 1.907$, $p = 0.167$, $\Delta AIC = -0.09$. The effect indicates that individuals have a lower chance to find the global optimum, if they receive a local optimum as a recommendation, $b = -0.634$, $SE = 0.058$, $OR = 0.531$, Wald $Z = -11.000$, $p < 0.001$. The absence of an interaction indicates, that this decrease is independent of whether the option was received as a recommendation or self selected in the corresponding click. However, the probability of receiving a locally optimal option was 14.73% if individuals received a recommendation, whereas it was 7.65% if they received none. Thus, a higher propensity of receiving a locally optimal recommendation than self-selecting one, could lead individuals to get stuck more likely with than without social information.

**How is distance to the recommendation affected?** Next, we ask whether recommendations attract exploration and therefore investigate the average distance to the recommendation. Figure 3.7A shows that individuals explore further away from low valued recommendations and closer to high-valued familiar recommendations across recipient experience and environment types. Thus, individuals might be attracted by high-valued familiar recommendations.

We further corroborate these observations with the results from a linear mixed model predicting the average distance to the recommendation after receiving it (or selecting its control option) on the round level. Results from this model are summarized in Figure 3.7B (for details on the model selection procedure see Appendix B.3). Estimates correspond to differences in the MD measured as number of tiles. Contrary to our expectations we do not find evidence for overall closer exploration to recommendations. Yet, we find evidence for a repulsion effect of low-valued recommendations.

Individuals explore closer to high-valued and familiar recommendations, as indicated

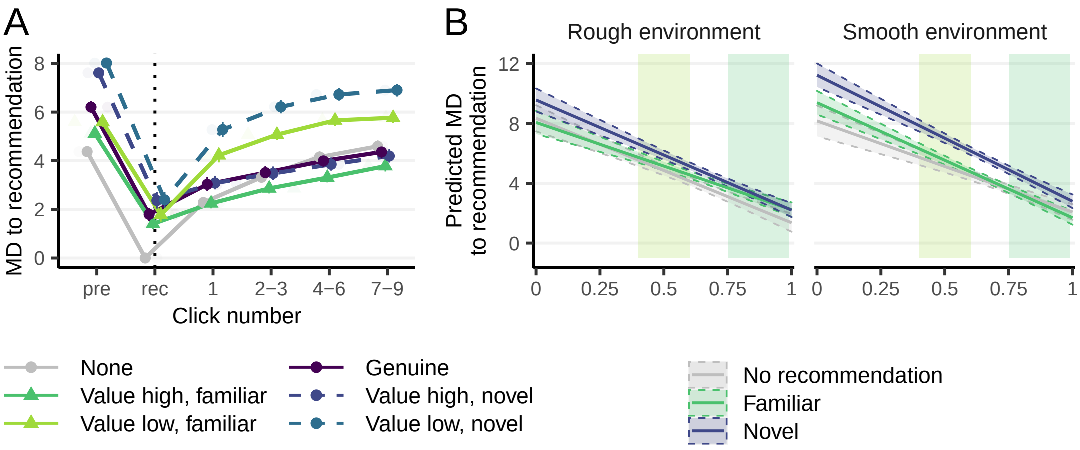

Figure 3.7: Average MD to the recommendation (A) descriptive over time (with click numbers collapsed in blocks) averaged across environment type and recipient experience and (B) estimated from a linear mixed model across time. Individuals explore closer to good recommendations and keep an increased distance to bad recommendations. The latter effect is especially pronounced among novel recommendations. Confidence bands correspond to 95% intervals. In (A) confidence intervals were cluster-bootstrapped on the individual level.Shaded regions in (B) indicate the lower and upper ranges of generated recommendation values respectively.

by negative main effects of recommendation value, $b = -0.367$, $SE = 0.016$, $t(684.68) = -23.314$, $p < 0.001$, and recommendation familiarity, $b = -0.467$, $SE = 0.052$, $t(421.19) = -9.045$, $p < 0.001$. This could reflect that individuals use the value of options in line with function learning accounts and follow localized exploration (e.g., Wu et al., 2018).

Moreover, results account for a repulsion from low-valued recommendations. A positive 2-way interaction of recommendation value and familiarity suggests that their joint negative effect on the distance to the recommendation is slightly weakened, $b = 0.028$, $SE = 0.013$, $t(2820.86) = 2.113$, $p = 0.035$, and the decrease of distance with recommendation value is stronger among novel than among familiar recommendations. Moreover, a positive interaction of recommendation value with environment type, $b = 0.037$, $SE = 0.014$, $t(3549.22) = 2.582$, $p = 0.010$, indicates that the attraction effect of recommendation value is weaker in rough environments. Also familiarity has less of an impact in rough environments, as indicated by a positive interaction between familiarity and environment type, $b = 0.213$, $SE = 0.048$, $t(3486.52) = 4.473$, $p < 0.001$. Put differently, individuals are more attracted by high-valued recommendations in smooth environments. Together these effects indicate, that individuals' exploration is driven away from novel, low valued recommendations, especially in smooth environments, reflecting a potential repulsion effect.

The repulsion effect also reflects in differences to self-selected options in the absence of a recommendation. A negative difference to no-recommendation trials, $b = -0.789$, $SE = 0.112$, $t(746.13) = -7.064$, $p < 0.001$, indicates that individuals explore closer to self-selected options than to average recommendations of intermediate value. Although a 2-way interaction of this difference with value is overall non-systematic, $b = 0.040$,

$SE = 0.028$, $t(2763.82) = 1.432$, $p = 0.152$, the relationship of value and receiving a recommendation varies with environment type, $b = -0.059$, $SE = 0.028$, $t(3156.19) = -2.092$, $p = 0.037$. This indicates that the decrease with recommendation value is larger in smooth environments, whereas it is more similar to the decrease without a recommendation in rough environments (see also Figure 3.7B).

Additional comparisons underscore this repulsion effect. Individuals explore further away from low valued (50 points) novel recommendations than from self-selected options of the same value both in smooth, $\Delta b = 0.369$, $SE = 0.093$, $t(1867.90) = 3.984$, $p < 0.001$, and rough environments, $\Delta b = 0.736$, $SE = 0.094$, $t(1915.03) = 7.797$, $p < 0.001$. In smooth environments, individuals also explore further from high valued recommendations (0.8), $\Delta b = 0.625$, $SE = 0.089$, $t(1713.18) = 7.048$, $p < 0.001$. Yet, the difference for low-valued options is larger, $\Delta b = 0.704$, $SE = 0.272$, $t(1942.02) = 2.590$, $p = 0.048$. Thus, individuals explore further away from novel low-valued recommendations than from self selected options and high-valued recommendations. This suggests that individuals avoid the areas around such recommendations (see also Figure 3.7).

Finally, positive interactions of recommendation value with the indicator for genuine recommendations, $b = 0.092$, $SE = 0.025$, $t(1634.76) = 3.711$, $p < 0.001$, and rough environment type, $b = 0.257$, $SE = 0.131$, $t(1222.20) = 1.965$, $p = 0.050$, as well as their corresponding positive 3-way interaction, $b = -0.055$, $SE = 0.026$, $t(1612.25) = -2.131$, $p = 0.033$, indicate that individuals in rough environments explore further from high-valued genuine recommendations. Consequently, individuals explore closer to genuine recommendations in smooth environments. This could, however, also be because genuine recommendations were on average larger, 0.79 ($SD = 0.24$), than other options, 0.66 ($SD = 0.18$). This restricted range could be related to a decreased slope.

In sum, how individuals are attracted by recommendations is modulated by multiple factors. Crucially, low-valued recommendations lead individuals to avoid the surrounding region, serving as red-flags. This indicates that individuals integrate the information conveyed by the recommendation. If we consider how individuals rely on recommendations after receiving them, they do not differ from self-selected options of the same value. This renders a function learning account with an always positive utility bonus less plausible as a description of exploration behavior over time. Hence, it is unlikely that recommendations are generally treated as intrinsically rewarding (Biele et al., 2009).

**Performance cost of continued exploration** Finally, we explored whether individuals change their overall exploration behavior after receiving a recommendation. The average MD to the previous option after a recommendation is (or would have been) received is 2.13 tiles. Moreover, we find that in 79.98% of all rounds individuals switch at least once before their final click. This indicates that most individuals do not completely exploit an option until the end.

The propensity to continue exploration instead of exploiting recommendations came at a cost. Individuals could have earned an average of 8.15 points ($SD = 5.98$) points more (range −4.55 to 32.08 points) by exploiting all recommendations better than their current best. This corresponds to an average loss of US-$ 0.14, had both rounds been rewarded. This would have been about 23.25% more than the earnings of US-$ 0.60 in an average round, in which a recommendation better than the current best option was received. Thus, it appears that exploiting received recommendations could have paid off more than what individuals actually did. Even if they fail to optimally exploit the available information individuals' strategies should not be conceived of as irrational, because they appear to robustly balance exploration and exploitation (see also Wu et al., 2018).

## 3.4 General Discussion

We investigated how the recommendation situation affects exploration behavior among advisors and recipients in spatially correlated bandits. Advisors changed their exploration behavior if their rewards depended on recipients' performance, but not if they faced the mere prospect of providing a recommendation. Recipients' exploration behavior is affected by clicking on recommendations, which is modulated by previous experience. They very likely clicked on recommendations. Thereby, they managed to avoid familiar low-value recommendations based on their previous experience in more predictable environments. Moreover, they used recommendations to inform their further exploration, exploring closer to good recommendations, while being repelled by bad recommendations. Although it would have often been superior, recipients rarely exploited good recommendations and rather continued to acquire further information themselves. In sum, the recommendation situation can affect exploration behavior of advisors and recipients, whereby recipients integrate recommendations with their previous information both prior to and after taking it.

### 3.4.1 How is exploration behavior affected?

Changed exploration behavior among advisors is hardly surprising, because it merely reflects that individuals maximize their rewards. Yet, it may have implications for real world situations. If advisors' rewards depend on their recipients (e.g., clients, Dana and Cain, 2015), they may explore more diverse information to make a recommendation (see also Benjamin & Budescu, 2015).

Recipients' exploration behavior is affected by a clear tendency to click on recommendations they receive. Crucially, how exploration behavior is affected depends on the familiarity and value of the recommendation. Thereby, individuals are more strongly affected by familiar recommendations. They more likely click on unknown familiar as compared to novel recommendations and explore closer to familiar recommendations. This likely reflects that individuals have a general propensity to explore close to their previous option. Within function learning this is captured by the idea of locality of exploration, which has previously been incorporated in function learning models in spatially correlated bandits (Wu et al., 2018). In fact, a preference for familiar recommendations is in line with previous research suggesting that distant advice (Harries et al., 2004) and distant opinions (Giese et al., 2019; Moussaïd et al., 2015) are discounted more strongly.

Locality of exploration or discounting of novel options could also explain why individuals are considerably less likely to click on recommendations immediately than to click on them later. Since individuals explore close to their previous option, novel options may have a lower likelihood of being clicked. However, the cumulative chance over time offsets this effect. Relatedly, this may indicate that they prefer to continue searching further information on their own as has been indicated by previous research (Schrah et al., 2006).

Importantly, locality resulting in a preference for familiar recommendations comes at a cost. In fact, individuals are more likely to find the best option in the environment if they receive a novel recommendation. This indicates that they may miss out on valuable novel information.

Yet, individuals' discounting of novel information is decreased if the environment is less predictable. The preference for familiar recommendations only seems to persist in more predictable, smooth environments. This is in line with social learning research, suggesting that individuals rely more on social information when the environment is uncertain or less predictable (Laland, 2004). Accordingly, the impact of recommendations is likely decreased, if individuals perceive the environment as more predictable.

Further, exploration is affected by the value of recommendations, even before the value can be known to the individual. Apparently, individuals generalize from their previous experience to inform their exploration behavior. Firstly, we find mixed evidence, partly in line with our expectation that individuals rely on their estimates of recommendation value in predictable environments. This is in line with our model which assigns a recommended option a bias. Yet, we do not find evidence for this effect, when considering if a recommendation is taken immediately. In fact, a general tendency to explore locally could have led individuals to initially dismiss many recommendations despite their increased value, since they first explored closer to already known options. Thus, whereas the judgment of a recommendation immediately may be less affected by its value, this is changed over the course of further exploration.

Beyond that, the value of recommendations also affects how individuals explore after a recommendation is received. Individuals explore closer to high valued recommendations and further away from lower valued recommendations. Thus, individuals behave in line with function learning, which estimates the value of unknown options to inform their probability of being selected.

Findings indicate that individuals incorporate recommendations into their new experience and generalize from it to inform their future choices. Potentially, they also use their previous experience to generalize to unknown recommendations. Thereby, a crucial question is whether recommendations are treated differently from self-selected options. It has been suggested that recommendations are intrinsically rewarding (Biele et al., 2009; Biele et al., 2011). Incorporating this idea, our model suggested that additional intrinsic rewards could lead individuals to get stuck, resulting in a decreased probability of find the best option. However, we only observe this decrease relative to self-selected options of the same value among high valued recommendations in smooth environments among the less experienced. Are recommendations then not intrinsically rewarding in the present setting?

A potential explanation is, that the additional intrinsic value of recommendations is limited to unknown options. In the task used by Biele and colleagues, the value of options varied and individuals could at best estimate their expected value (Biele et al., 2009; Biele et al., 2011). Likewise, in the present study, a high probability to click on recommendations suggests that they are treated differently, as long as they are unknown. Yet, after the recommendation was revealed, its uncertainty was strongly decreased. Thus, again in line with research on social learning, individuals rely more on social information when the environment is uncertain or less predictable (Laland, 2004).

### 3.4.2 Limitations

Notwithstanding, the present results are subject to some limitations. Results relying on environment type face the issue that the type of environment was randomized rather than manipulated. Thus, one can argue, that a causal interpretation is not warranted here. Yet, we provide theoretically sound explanations for the effects found that rather warrant replication by future studies.

Moreover, we included an individual control, as suggested by recent reviews (Miton & Charbonneau, 2018). The control conditions without recommendations included as many clicks as the conditions with recommendations. Yet, they do not allow to compare the full amount of additional information inherent in recommendations (which were an additional 20 clicks among advisors). Thus, the comparison may serve less well to draw conclusions about advantages for the performance of dyads relative to the effort put into information search, since in sum there was more information search involved when a recommendation was received. Nevertheless, we argue that this comparison still serves well to answer how individuals are affected by the recommendation as a single additional piece of information.

Finally, conditions for more maladaptive transmission dynamics may largely not have been met. Individuals rely more on social information if individual trial-and-error learning is costly (e.g., Morgan et al., 2012). Yet, here individual information was comparably cheap, because individuals had 5 or 15 clicks before and 10 clicks after they received a recommendation. Similarly, the only indication of a disadvantage of receiving a recommendation we find among the least experienced. Crucially, a shorter time horizon could render individual information more costly (see Berger-Tal et al., 2014) and lead individuals to rely more strongly on recommendations. Future work should also address this possibility. Despite these limitations we argue that our findings can inform future work on the transmission of recommendations beyond dyadic interactions.

### 3.4.3 Future directions: Predicting the transmission of recommendations

The above influences need to be considered when addressing the question whether and how bad or inferior recommendations can persist over time. In transmission chains multiple individuals perform a task and pass on information or solutions to another individual (e.g., Mesoudi & Whiten, 2008). The present results can inform the predictions of such chains, which pass on recommendations selected by individuals instead of last solutions (e.g., Caldwell & Millen, 2009; Yahosseini & Moussaïd, 2020) or solutions from individuals who explore the environment concurrently (e.g., Acerbi et al., 2016; Mesoudi & O'Brien, 2008a).

A first question is to what extent the behavior of recipients can be generalized to a situation in which they pass on a recommendation themselves. Importantly, we find that advisors do not change their behavior in a recommendation situation, in which their rewards do not depend on the performance of the recipients of their recommendation. This indicates that one could extrapolate the present results in simulations before investigating the predictions experimentally.

A second question is by which factors the performance towards finding the best solution is affected. Our results indicate that transmission chains will especially benefit from the transmission of novel information. This is reminiscent of the finding that the Wisdom of the Crowds (Surowiecki, 2004) benefits from the aggregation of uncorrelated information (Lorenz et al., 2011).

In the present setting, individuals tend to continue exploration, even if they receive good recommendations they could exploit. Thus, if individuals can acquire individual information before and after receiving a recommendation, they will likely contribute novel information. Moreover, generalization can allow them to filter out bad recommendations in advance. Together, this will likely facilitate the progress towards the best alternative. Yet, these benefits likely disappear, as soon as individuals have less capacity for individual exploration (before and after receiving a recommendation). If the strong reliance on familiar recommendations persists, this may slow down the progress towards the best option. Moreover, decreased information from which to generalize may take away the possibility to filter out bad recommendations.

Future work should investigate these influences in the context of transmission chains that move beyond a single dyadic link. In order to get more precise predictions one could extrapolate the present results in simulations and investigate the predictions experimentally. Thereby, it would be important to vary individuals' degree of experience. This way our results from a dyadic link could fruitfully foster future work on the transmission of information and cultural evolution.

### 3.4.4 Conclusion

Relying on recommendations and social information is part of our human nature, as is generalizing from previous experience. Nevertheless, we often do not rely on information provided by others and fail to tell bad from good advice. Previous experience can inform the judgment of recommendations and recommendations can inform further individual learning. We were the first to investigate how the prospect of providing a recommendation, as well as receiving a recommendation, affect information search in spatially correlated bandits. In the light of our results humans are adapted well to benefiting from recommendations to identify better options, yet sometimes could rely more heavily on them. However, a persistent tendency to continue exploration makes it likely that individuals will be able to pass on increasingly better options to others which is essential for successful cultural transmission and avoidance of maladaptive informational cascades. Improving our understanding of such low-level transmission and integration processes like in advice giving and taking, will eventually also help to better understand collective functioning and human cultural evolution.

# 4

# Allocation of scarce resources

## Moral reasoning in the wild: Triage decisions in the context of the COVID-19 pandemic reflect integration of information beyond a concern for the greater good

*Joint work with Hansjörg Neth.*[3,4] *and Wolfgang Gaissmaier,*[1,2,3,4]

[1] *Graduate School of Decision Sciences, Universität Konstanz*
[2] *Graduate School of the Social and Behavioural Sciences, Universität Konstanz*
[3] *Department of Psychology, Universität Konstanz*
[4] *Centre for the Advanced Study of Collective Behaviour, Universität Konstanz*

## Abstract

Triage decisions about who receives scarce, life-saving treatment during a pandemic offer a window for studying factors that affect moral reasoning in an ecologically valid scenario. Moral reasoning is often in line with normative utilitarian considerations but is modulated by multiple factors beyond that. Although triage decisions strongly rely on patient features related to saving a larger number of individuals (like patient survival probability) additional features are known to play a role. It is an open question which patient features are used in allocation decisions and how they interact with with characteristics of the decision maker. In an online study, 1753 participants each evaluated 19 moral dilemmas that required allocating the only available ventilator to save one of two patients, who differed in one focal feature. In 2 of 4 collection waves we added a default condition that manipulated whether the ventilator would be withdrawn from a patient. Our results indicate that the patient whose survival conveys the larger (group) utility receives the ventilator more likely across features related to saving a larger number of individuals and features which are not. Preferences according to features can further be explained by shared group-membership between participant and patient and could be related to altruistic (third-party) punishment. Participant characteristics like religiosity, political orientation, and perceptions of the likelihood and severity of a COVID-19 infection have an impact on allocation decisions. Crucially, across many features the ventilator is more likely withdrawn from the patient who conveys a lower level of group-level utility. Our results demonstrate that various pieces of information are integrated in hypothetical triage decisions.

## 4.1 Introduction

Scarcity of life-saving resources like ventilators in the COVID-19 pandemic made triage decisions, which require a trade-off whose life is saved, disturbingly real (Hick et al., 2020; Solnica et al., 2020; White & Lo, 2020). People's responses to such realistic triage situations in a pandemic can provide valuable insights into what shapes moral reasoning, because the traditional trolley problems (Foot, 1967; Thomson, 1985) are artificial (Bauman et al., 2014; Bostyn et al., 2018) and likely miss out on important contextual factors (Camerer & Mobbs, 2017; Carnes et al., 2015; FeldmanHall et al., 2012). Utilitarian considerations (Bentham, 1983; Mill, 1863; Singer, 2011 cf. Kahane et al., 2018) prescribe that triage decisions should be made with the goal of creating the largest benefit for the largest number of individuals (Persad et al., 2009), irrespective of their group membership (see e.g., Kahane et al., 2018). Accordingly, ethical guidelines recommend saving those with the better prognosis or relevant to the healthcare system in order to save a larger number of people (Emanuel et al., 2020; Jaziri & Alnahdi, 2020; Savulescu et al., 2020). Laypeople largely endorse such utilitarian policies (Buckwalter & Peterson, 2020) and approve of allocating care to those with the better prognosis (Fallucchi et al., 2020) but also rely on features that are irrelevant to the number of people saved (Jin et al., 2021; Wilkinson et al., 2020). Thereby, it is an open question which factors of the decision maker, the decision situation and the real-world context affect whether a patient feature is used in triage decisions or not.

Laypeople's triage decisions largely align with guidelines, as well as further cost-benefit considerations (Fallucchi et al., 2020; Wilkinson et al., 2020). Beyond survival probability and time in the ICU they are strongly affected by age, and whether a patient has dependent children (Wilkinson et al., 2020). For both one can argue that the benefit of saving the younger person or a person with dependents is larger because more life-years are saved and at least one additional person benefits respectively (also age is a risk factor for severe COVID infections). Likewise, people preferably allocate scarce medical resources to patients like health professionals who can help saving more people during the crisis (Fallucchi et al., 2020).

Although moral reasoning is often in line with cost benefit calculations (Conway et al., 2018), common sense moral reasoning does not strictly follow utilitarian cost-benefit calculations (Everett & Kahane, 2020; Holyoak & Powell, 2016; Waldmann et al., 2012). Instead, it has been argued that moral decisions are based on intuitions (Cushman et al., 2006; Haidt, 2001) and emotions (Greene & Haidt, 2002), a universal moral grammar (Mikhail, 2007) or simple heuristics (Gigerenzer, 2010; Sunstein, 2005). Moral judgments can be utilitarian in a more or less strict sense, from involving any kind of cost-benefit calculations ignorant of the greater good, over cost-benefit calculations aiming at the greater good, to reasoning explicitly with impartial utilitarian principles (Conway et al., 2018).

A multitude of factors beyond saving the larger number of people affects moral reasoning (Christensen & Gomila, 2012; Rai & Fiske, 2011; Van Bavel et al., 2015). First, moral decisions are not necessarily impartial (Kahane et al., 2018). People treat those favorably with whom they share a group (Everett et al., 2015; Tajfel et al., 1971; Yamagishi & Kiyonari, 2000). This could be driven by values like identity and reciprocity that prefer those from one's group (Rai & Fiske, 2011). Accordingly, people are less likely to save people from extreme out-groups in trolley dilemmas (Cikara et al., 2010) and more willing to accept harm against foreigners (Uhlmann et al., 2009). In triage decisions this reflects in a preference for saving those of a shared nationality, which is more pronounced among conservatives (Jin et al., 2021). It is, however, unclear whether partiality for one's in-group in triage decisions generalizes to other groups, like people who recently migrated to a country, self-selected groups like religiosity, or non-self selected groups like ethnicity.

Generally, saving more cooperative individuals who contribute to society benefits the group by increasing the likelihood of cooperation and thus may positively enter cost-benefit calculations. People maximize within-group cooperation by punishing non-cooperative individuals at their own cost (de Quervain et al., 2004; Fehr & Gächter, 2002; Fowler, 2005), even if they are themselves unaffected (Balliet et al., 2011; Fehr & Fischbacher, 2004; FeldmanHall et al., 2014; Henrich et al., 2010) and will not interact with the other in the future (Boyd et al., 2003). The tendency for punishing non-cooperators is especially pronounced among individuals who cooperate themselves (Fehr & Fischbacher, 2004). Because most individuals in the U.S. (80–90%) consider paying their taxes as an important duty (IRS, 2019), tax evasion may elicit such punishment. Also more contentious cooperative behaviors like vaccination (Korn et al., 2020) and organ donation (Dijker et al., 2013; Stijnen & Dijker, 2011) could elicit a preference that is more pronounced among those who show the cooperative behavior. It is, however, unclear whether these features also elicit punishment in life or death decisions.

Other factors that could affect triage decisions are the severity of the crisis and perceptions thereof. For moral decisions, not only features of the decision targets but also context matters (Carnes et al., 2015; Decety & Cacioppo, 2012; FeldmanHall et al., 2012). Dependent on how realistic and self-relevant the decision situation is perceived, people could become more or less likely to save those with higher survival probability and instrumental value. Specifically, utilitarian considerations are endorsed more in the context of COVID-19 than in a generic context (Kneer & Hannikainen, 2021; Navajas et al., 2021). Results from triage decision suggest that endorsement of survival probability increases in the actual severity of the crisis, whereas it decreases in perceiving the crisis as more severe (Jin et al., 2021). We add to this research and ask how case numbers and perceptions like having had an infection, perceived infection severity and likelihood, as well as the severity of the outbreak affect the decision how life-saving resources should be allocated.

Beyond the factors discussed so far, it is very likely that decision makers are also sensitive to whether care has to be withdrawn from an individual and who that individual is. Generally, people avoid directly harming others in trolley dilemmas (Cushman et al., 2006; Foot, 1967; Thomson, 1985; Waldmann et al., 2012) and avoid responsibility for aversive decision outcomes (Leonhardt et al., 2011). Accordingly, people have a strong preference for first come, first serve allocations in triage decisions (Fallucchi et al., 2020), despite not being medically relevant (Biddison et al., 2014). Still they are willing to re-allocate the ventilator to a patient with a better prognosis (Wilkinson et al., 2020). We address the question whether their preference for first come, first serve or inaction eliminates the effects of patient attributes unrelated to survival or whether they are more likely to withdraw the ventilator from some patients than from others

In fact, people may also disguise their preference for re-allocation by allocating the ventilator randomly. Although people dislike uncertainty (Simonsohn, 2009), and random allocation in particular (Keren & Teigen, 2010), people also have a strong preference to appear fair (Andreoni & Bernheim, 2009). People use randomness to avoid responsibility for decision outcomes (Dwenger et al., 2012) and disguise their self-interest behind uncertainty or randomness (Batson et al., 1999; Dana et al., 2007). Moreover, implicit attitudes affect what individuals perceive as fair. For instance, people may punish out-group members by rejecting unfair offers more than members of their in-group (Kubota et al., 2013). Decision makers could balance their goal to appear fair with their preference to allocate the ventilator to their favored patient by choosing random allocation if the ventilator would have to be withdrawn from the disfavored patient. This would reflect in more random allocations if the ventilator has to be withdrawn from the disfavored than from a favored patient.

We present results from a cross-sectional online study in which U.S.-based participants indicated to which of two patients who differ in a single feature the only available ventilator should be allocated to save one patient's life. We distinguish patient features of two types: The first type is potentially relevant to promote the greater good. Features of this type are risk factors related to patient's survival, more life years saved, or provides additional benefits through affecting dependents. The second type is not clearly related to promoting the greater good and captures various other factors like patients' group membership and contributions to society. In 2 of 4 waves we also experimentally manipulated the default allocation. That is, whether the ventilator was already allocated to one of the patients so that it has to be withdrawn. In contrast to previous studies which investigated a larger number of features (Awad et al., 2018; Jin et al., 2021) participants could opt to allocate the ventilator randomly. Not only is random allocation the normative response to dilemmas in which patients do not differ in relevant medical characteristics (Emanuel et al., 2020; Jaziri & Alnahdi, 2020; Savulescu et al., 2020). It is also crucial to address the question, whether people are actually willing to consider certain patient features or whether they just use any discriminating information they are given to make a decision. Beyond our manipulations we collected participant characteristics, including demographics, and perceptions of the current COVID-19 pandemic. Counties of residence additionally allowed us to assess the number of infections in the participant's county of residence.

## 4.2 Methods

### 4.2.1 Participants

In total, we recruited 1868 participants on Amazon Mechanical Turk (40.9% female, mean age 37.7 years, range 18 to 83; see Table C.3 for a full sample description). We collected data in 1-week intervals between March 2nd and May 23rd.[1] Data were always collected on Thursday evening (starting at around 5 p.m. German time; UTC+1) so that we always reached Mturk workers during the day in the US (beginning from about 8 a.m. to about 2 p.m and lasting to about no later than 5 p.m.) in order to avoid possible effects of time of day, weekday on the participant pool (Arechar et al., 2017). Participants provided informed consent prior to participating in the study. Checks of entry plausibility and completion time, as well as self-reports of data quality were used to screen out participants who did not pay sufficient attention. We excluded 77 participants, resulting in a total $N$ of 1753 (about 225 per condition/wave).

### 4.2.2 Design

Each participant made a total of 19 hypothetical decisions about to which of two patients a doctor should allocate the only available ventilator during a viral epidemic that had created a scarcity of medical resources. The two patients in each of the scenarios were described to differ in exactly one focal feature. Participants could either opt to deliber-

---

[1] We also collected data on August 27 and September 3 but the MTurk participant pool changed substantively (Arechar & Rand, 2020), so that it is hardly possible to disentangle effects of the COVID-19 crisis, political events and differences in the participant pool. Therefore, the present publication is limited to waves 1 to 4.

ately allocate the ventilator to one of the two patients or to use a fair, random device (like a lottery or coin). This allowed us to address the question how different features affect allocation of the ventilator. Based on our expectations who will be treated preferentially due to a higher perceived utility (due to more lives saved or higher contributions to society), we speak of favored and disfavored patients in the following. For a list of patient pairs see Figure 4.1 (pairs are aligned, so that the individual we expected to be saved more likely is placed on the right hand side). It was made clear that the patient who received the ventilator would be saved, whereas the other patient would die, thus eliminating uncertainty about the outcome.

In waves 3 and 4 we additionally varied between-subjects whether the favored patient, the disfavored patient, or neither already had received the only available ventilator before. Thus, the decision became whether to keep the ventilator assigned to the default individual, withdraw and reallocate it or decide about re-allocation by means of a random device. An example scenario is provided in the supplement, Figure C.1.

### 4.2.3 Procedure

First, participants provided informed consent and received instructions that they will be asked to make decisions about who of two comparable patients differing in one key regard should receive life-saving treatment. Patient pairs were presented in a random order for each participant. Next, participants answered questions about their risk and severity perceptions of the current Corona crisis, their own health, altruism, time and risk preferences, as well as demographic questions, including their county, their political views, and their religiosity. To allow for the detection of in-group effects we asked participants to provide information on characteristics equal or related to the characteristics of our hypothetical patients (e.g., health, alcohol consumption, family and relationship status). Moreover, they provided information on their county so that we were able to assess the objective severity of the COVID crisis in their area of residence (an overview over participant demographics can be found in Appendix C.3). Participants always had the option to leave questions blank.

### 4.2.4 Measures

To assess participants' perceptions of the crisis we used 7-point scales translated from (Betsch et al., 2020) asking for the probability (1: extremely unlikely to 7: extremely likely) and severity of an infection (1: extremely harmless to 7: extremely severe), as well as their perceptions how severe the outbreak was in their area of residence (1: completely harmless to 7: very severe). Questions about time- and risk-preferences, as well as altruistic giving were simple 1-item measures taken from the preference survey module (Falk et al., 2016; Falk & Hermle, 2018). To assess altruism participants indicated how much of US-\$1,000 they are willing to spend to a good cause. Political views and religiosity were assessed with simple items using a 7-point scale from conservative to liberal and non-religious to religious as used in Awad et al. (2018). We asked about alcohol consumption as standard drinks per day (Babor et al., 2001). In our analysis we contrasted moderate alcohol consumption (0 to 2 standard drinks per day) with hazardous alcohol consumption (more than 2 standard drinks per day). To assess whether participants were overweight, we asked them to provide their height and weight. We excluded participants with implausible entries. Finally, we asked participants whether they are registered as an organ donor, whether they got the flu shot in the current season and about their attitude towards flu vaccination is on a 5-point scale (1: strongly oppose, 5: strongly favor). Also other questions were included, which are not part of this publication, which is part of a larger project.

### 4.2.5 Data analyses

As our main dependent measure, we report the probability that the favored patient is assigned the only ventilator (across participants). This probability makes use of all available information about allocations. To account for non-independence between allocation decisions, we use binomial generalized mixed modelling. A binomial generalized model predicts the number of successes on a given item. Because we are interested in the probability that the favored patient receives the ventilator we count allocation to the favored patient as 2 of 2 successes, random allocation as 1 of 2 (because the favored individual can be expected to receive the ventilator in 1 out of 2 cases), and allocations to the disfavored patient as 0 of 2 possible successes. This way the model estimates the probability that the favored patient receives the ventilator and survives. Additionally, we use logistic models predicting the probability of random allocations. We control for participant and item variation by including random intercepts for both (Baayen et al., 2008). For model tables presenting coefficient estimates see Appendix C.4.

## 4.3 Results and Discussion

From participants' responses we can estimate the expected probability that the favored patient receives the ventilator assuming that a random participant makes the allocation decision (Allocation to favored: $p_{favored} = 1$, random allocation: $p_{favored} = 0.5$, allocation to disfavored: $p_{disfavored} = 0$). The larger the proportion of participants who allocate the ventilator to the favored patient is, the higher this patient's probability of receiving the ventilator and surviving would be. Random allocations draw the probability closer to chance. We analyze this probability with a generalized binomial mixed model with random intercepts for participants and features to account for cross-classification of observations. To address whether participants select the normative response on features unrelated to the greater good, we additionally use a generalized logistic mixed model on the probability of a random allocation.

We first consider how features and their interaction with participant characteristics affect from whom the ventilator is withheld (no default). Thereby, we also address whether participants are more likely to select random allocation among features unrelated to the greater good. We do not find effects of time-point of data collection on the withholding of care, $\chi^2(3) = 3.543$, $p = 0.315$. Therefore, we show results across all 4 waves. Finally, we address how changing the context from withholding to withdrawing care (experimental manipulation in time-points 3 and 4) modulates the feature effects.

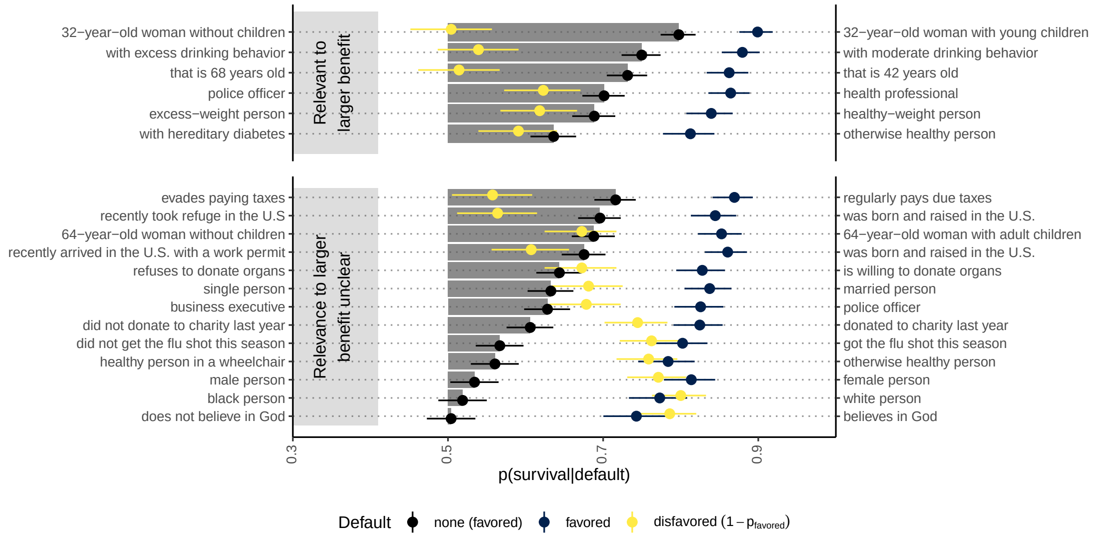

Figure 4.1: Estimated expected probability that a patient receives the ventilator, aggregated over directed and random allocations. The left and right labels correspond to the features of the patients we expected to be disfavored and favored respectively. Features are grouped by whether they are clearly relevant to considerations of the greater good or not. Grey bars and blue points indicate the probability of the favored patient to receive the ventilator if no patient has it (withholding). For the respective default-conditions, grey points show the probability that the favored default can keep the ventilator and yellow points show the probability that the disfavored default can keep the ventilator (1 minus the probability that the ventilator is re-allocated to the favored). Probabilities differ by feature. Overall, the patient who has the ventilator has a probability of keeping it that is higher than the probability of receiving it in situations of withholding. Error bars are 95% confidence intervals estimated from a generalized binomial model.

The grey bars and black points in Figure 4.1 show the probability of survival of the favored patient across features. Cost-benefit calculations most clearly apply to features where the relevance to the greater good through benefiting more individuals (e.g., health-related features, young mother) or yielding a larger benefit (age) is clear. The favored patient receives the ventilator more likely across features relevant to the greater good, 71.1%, $CI_{95\%} = [66.4; 75.4]$, than among other features, $OR = 1.58$, $CI_{95\%} = [1.22; 2.05]$, $\chi^2(1) = 11.905$, $p = 0.001$. Yet, the favored patient also receives the ventilator more likely than than the disfavored patient among features where the relation to the greater good is unclear, 60.9%, $CI_{95\%} = [57.3; 64.4]$. Results from a model with fixed effects for each feature reveal that the only features for which the overall bias is not systematically different from 50% are ethnicity, 51.9%, $CI_{95\%} = [48.8; 55.0]$, and religiosity, 50.4%, $CI_{95\%} = [47.3; 53.5]$.

Importantly, according to triage guidelines the normative response in the absence of features implying different prognosis is random allocation. In fact, participants selected random allocation more likely if patients differed by features *unrelated* to utilitarian benefits (66.6%, $CI_{95\%} = [64.1; 69.2]$) whereas they selected random allocation less likely if patients differed by features *related* to utilitarian benefits (33.3%, $CI_{95\%} = [30.4; 36.2]$), $b = 1.508$, $SE = 0.044$, $OR = 4.519$, Wald $Z = 34.111$, $p < 0.001$, see Figure 4.2. Thus the strength of bias and the probabilities of random allocation concur with previous re-

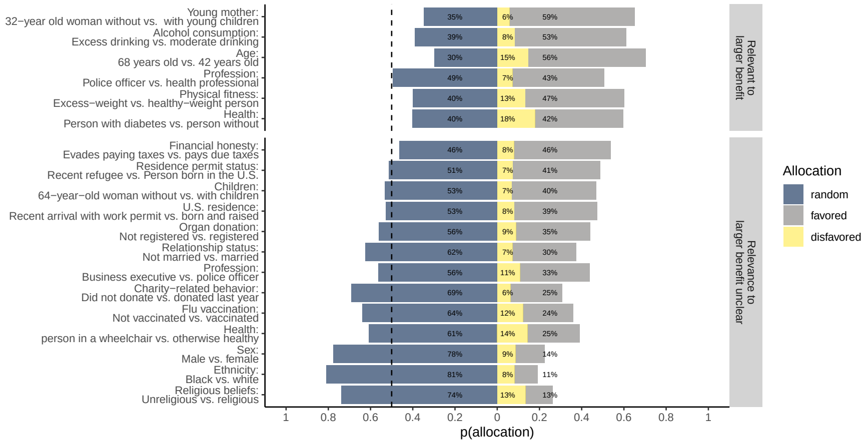

Figure 4.2: Distribution of allocations for each feature in situations of withholding care (no default). We grouped features by whether they are clearly relevant to considerations of the greater good or not. Among most features not relevant to the greater good a majority of participants selects the normative response of random allocation.

search suggesting that withholding care is strongly affected by features which are in line with utilitarian considerations (Fallucchi et al., 2020; Jin et al., 2021; Wilkinson et al., 2020). People appear to distinguish features also in the absence of given survival probabilities. However, several other features also affect participants' preferences affecting patients' likelihood to receive the ventilator. This preference is driven by participants who prefer to allocate the ventilator to the favored patient.

### 4.3.1 Are preferences affected by the COVID-19 crisis?

We expected that the severity of the crisis and perceptions thereof affect moral reasoning. Overall, instead of actual occurrence of infections, rather perceptions of and personal concern with the crisis affect triage decisions. Because case numbers, perceptions of and concern with the crisis were measured rather than manipulated, the following results are controlled for demographic variation. The perceived severity of the COVID-outbreak is only weakly positively related to cumulative incidence, $r = 0.186$, $p < 0.001$, whereas perceived likelihood, $r = 0.039$, $p = 0.094$, and severity, $r = 0.056$, $p = 0.016$, show a systematic but negligibly small relationship to cumulative incidence. Because perceptions of the crisis are largely unrelated to cumulative incidence, we can assess their independent contributions to moral reasoning.

Cumulative incidence in participant's county of residence is not systematically related allocation probabilities, $\chi^2(1) = 0.463$, $p = 0.496$. Perceptions of the infection as more severe are related to a decreased probability of allocating the ventilator to the favored individual. Specifically, if individuals perceive an infection as *more severe* the effect is driven by *features relevant to the greater good*, $b = -0.082$, $CI_{95\%} = [-0.144; -0.019]$, $OR = 0.922$, and not by other features, $b = 0.016$, $CI_{95\%} = [-0.033; 0.065]$, $OR = 1.017$. Yet, if individuals perceive an infection as *more likely* the effect is driven by *features not relevant to the greater good*, $b = -0.066$, $CI_{95\%} = [-0.118; -0.015]$, $OR = 0.936$ and not by other features, $b = 0.011$, $CI_{95\%} = [-0.054; 0.077]$, $OR = 1.011$. Among those who report having had an infection themselves the favored patient receives the ventilator less likely, $\chi^2(3) = 13.023$, $p = 0.005$. The difference is systematic for features related to the greater good, $OR = 0.70$, $CI_{95\%} = [0.49; 1.00]$, but not for the other features, $OR = 0.83$, $CI_{95\%} = [0.62; 1.11]$. Finally, the probability that a health professional rather than a police officer receives the ventilator is positively related to how severe participants perceive a COVID-infection, $b = 0.086$, $SE = 0.033$, $OR = 1.090$, Wald $Z = 2.586$, $p = 0.010$. Importantly, there is no evidence for an increase (if anything, evidence favors a decrease) in random allocations across these effects (all $OR$ between $0.73, 1.04$). This suggests that individuals increase their deliberate allocations to the respective individuals and not their random allocations.

Together, this adds to the finding that the number of cases is positively related to the use of survival probability whereas crisis perceptions are not (Jin et al., 2021). Personal concern and taking an infection more seriously could be related to stress increasing altruistic tendencies (von Dawans et al., 2012), resulting in saving those more who are worse off (instead of those who survive more likely) and those who help others survive. Conversely, perceiving an infection as more likely could be related to increased self-relevance and alters utilitarian calculations based on other features. In part these results could also reflect benevolent partiality (Paolacci & Yalcin, 2020) in favor of those who one expects to be worse off.

### 4.3.2 Do we observe in-group favoritism and altruistic punishment?

Allocations to the favored individual differ predictably by participant characteristics, see Figure 4.1. A person born and raised in the U.S. receives the ventilator more likely than a person who came to the U.S. with a work permit, 67.6%, $CI_{95\%} = [64.7; 70.3]$, or a refugee, 69.6%, $CI_{95\%} = [66.8; 72.3]$. A more liberal political orientation is negatively related to the probability that the favored patient receives the ventilator across features, $b = -0.192$, $SE = 0.028$, $OR = 0.825$, Wald $Z = -6.929$, $p < 0.001$. This effect is strongly driven by fewer allocations to patients born in the U.S., $\chi^2(1) = 48.512$, $p < 0.001$, $b = -0.141$, $SE = 0.020$, $OR = 0.868$, Wald $Z = -6.965$, $p < 0.001$. This extends the finding that more conservative individuals base triage decisions more on nationality (Jin et al., 2021). Overall, individuals prefer to allocate the ventilator randomly if patients differ in their religiosity. However, higher self-reported religiosity is positively related to the probability that the favored patient receives the ventilator across features, $b = 0.084$, $SE = 0.032$, $OR = 1.088$, Wald $Z = 2.633$, $p = 0.008$. This effect is stronger if patients differ in their religiosity, $b = 0.165$, $SE = 0.025$, $OR = 1.179$, Wald $Z = 6.484$, $p < 0.001$. Evidence for an effect of participant ethnicity, $\chi^2(2) = 7.568$, $p = 0.023$, indicates that, although inconclusive, white participants tend to allocate the ventilator more likely to a white patient (or vice versa), $OR = 1.32$, $CI_{95\%} = [0.95; 1.82]$, see also Figure 4.3. In sum, group membership affects allocation decisions. Decision makers prefer to save people born in their country and people who share their worldviews. This underscores previous findings that religiosity and a conservative political orientation affect consequentialist thinking (Piazza & Sousa, 2014). Generally, decision makers' values strongly affect moral behavior (Rai & Fiske, 2011), not only directly but also through sharing certain worldviews or moral convictions (Bilancini et al., 2020; Parker

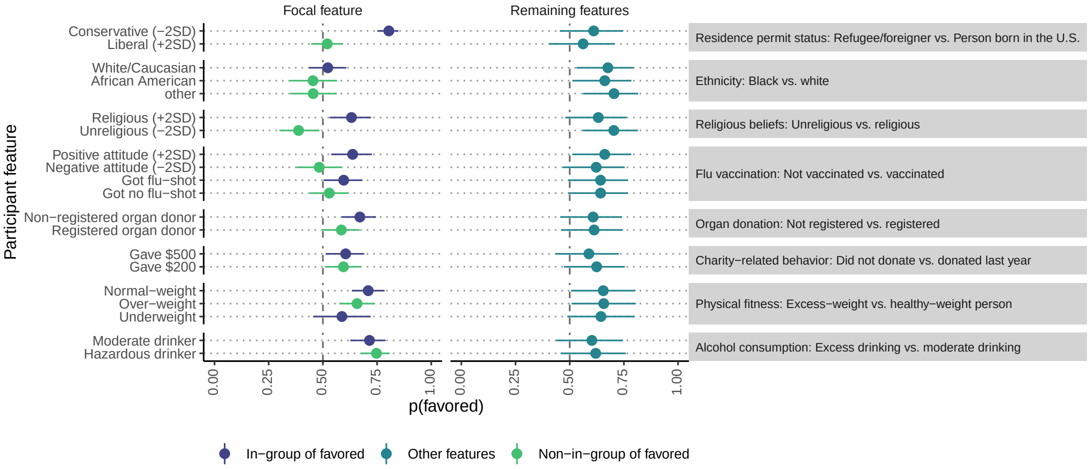

Figure 4.3: Probability that the favored patient receives the ventilator among participants whose characteristics indicated shared group membership with the favored patient or not (left panel). The right panel shows the effect of participants' characteristics across other features. Shared group membership (blue points) is related to a larger probability that the respective patient receives the ventilator than among those who do not share group membership (green points).

& Janoff-Bulman, 2013).

We expected that people prefer cooperative patients who contribute to society. In line with this expectation, the tax payer, the person vaccinated against the flu, the person willing to donate organs, and the charitable person receive the ventilator with a higher probability than their counterparts who do not show the respective behavior, see Figure 4.1. Moreover, we provide evidence that those who show cooperative behavior themselves show this favoritism more strongly (Fehr & Fischbacher, 2004), see Figure 4.3. Accordingly, the strong effect of tax evasion is potentially due to the widespread approval of paying taxes (IRS, 2019). Additionally, the probability that a patient who is vaccinated against the flu receives the ventilator is positively related to a more positive attitude to flu vaccination $\chi^2(1) = 3.946$, $p = 0.047$, and is be higher among those who report being vaccinated against the flu themselves, $\chi^2(1) = 4.858$, $p = 0.028$. Moreover, the probability that a patient who is registered as an organ donor receives the ventilator is increased among participants who report being registered as organ donors themselves, $\chi^2(1) = 10.990$, $p = 0.001$. However, we do not find evidence for a positive relationship of hypothetical donations and allocations to an altruistic individual, $b = 0.017$, $SE = 0.011$, $OR = 1.017$, $\chi^2(1) = 2.396$, $p = 0.122$. We also do not find substantive interactions between whether participants were normal-weight (according to BMI) and survival of the normal-weight patient, $\chi^2(3) = 7.502$, $p = 0.058$, or between participants with moderate drinking behavior (according to the AUDIT) being related to more likely survival of the moderately drinking patient, $\chi^2(1) = 0.465$, $p = 0.495$.

### 4.3.3 How are allocations changed in contexts of withdrawal?

Finally, one would expect that people avoid direct harm (Thomson, 1985; Waldmann et al., 2012) and are therefore less willing to withdraw the ventilator from either patient. Our experimental manipulation in waves 3 and 4 allows comparisons between whether and from whom the ventilator would have to be withdrawn. In fact, survival probabilities suggest that a patient who has been allocated the ventilator previously would keep it with a probability of 73.97%, $CI_{95\%} = [72.19; 75.75]$, see blue and yellow points in Figure 4.1. However, patient features retain a substantive impact on the decision to withdraw the ventilator. First, from whom the ventilator has to be withdrawn matters. Overall, participants were more willing to withdraw the ventilator from a disfavored patient. The favored patient keeps the ventilator with a probability of 84.34%, $CI_{95\%} = [81.75; 86.94]$, whereas the disfavored patient keeps it with a probability of 63.60%, $CI_{95\%} = [59.10; 68.09]$. Contrary to our expectations that people may disguise their preferences to appear fair (Anderson, 2003), we do not observe more random allocations if the ventilator has to be withdrawn from a disfavored as compared to a favored patient, $OR = 0.93$, $CI_{95\%} = [0.65; 1.34]$. Second, the asymmetry differs heavily by feature. The effect of who has the ventilator is increased among features not relevant to the greater good, $\chi^2(2) = 25.264$, $p < 0.001$. Among these features the favored patient keeps the ventilator more likely, $b = 0.080$, $SE = 0.017$, $OR = 1.083$, Wald $Z = 4.783$, $p < 0.001$, and receives it less likely if the ventilator would have been withdrawn from the disfavored patient, $b = -0.069$, $SE = 0.016$, $OR = 0.933$, Wald $Z = -4.350$, $p < 0.001$. The withdrawal situation has more impact among features not relevant to the greater good. This reflects that features relevant to the greater good have a stronger impact beyond the situational variable whether the ventilator has to be withdrawn from a patient.

## 4.4 General discussion

It is an open question how the interaction of patient features with characteristics of the decision maker and the decision situation affects allocation decisions in triage situations. Even in the absence of numeric survival probabilities, decision makers are more likely to save those whose survival benefits a larger number of people or yields a larger overall benefit. As predicted, features related to patients' contributions to society and social usefulness affect who is more likely to receive life-saving treatment. However, a majority of individuals prefers random allocation among these features. Contrary to expectations, perceptions of infection severity and personal concern weaken utilitarian tendencies except for saving health professionals who have instrumental value for solving the crisis. Interactions of decision maker's and patient's group membership underscore that the perceived value of a patient depends on who is asked. Thereby, most biases persist beyond situations of withholding care to situations in which instrumental harm means active withdrawal of care. This is in line with moral reasoning strategies that integrate information and are sensitive to cost-benefit analyses. Further research is required to understand the precise processes underlying the integration of these various pieces of information.

The strong endorsement of survival-related patient features is in line with utilitarian cost-benefit calculations (Fallucchi et al., 2020; Jin et al., 2021; Wilkinson et al., 2020).

Yet, the finding that other features affect allocation decisions as well indicates that patient utility enters moral reasoning beyond saving the larger number of individuals. Preferences for saving those with whom one shares a group membership and the more cooperative reflect that the utility for in-group individuals and cooperative individual may be increased. This aligns with findings on motivated moral reasoning (Parker & Janoff-Bulman, 2013) and in-group favoritism (Everett et al., 2015; Tajfel et al., 1971), as well as third-party punishment (Fehr & Gächter, 2002). Although the present data do not allow us to conclusively disentangle the role of third-party punishment and in-group favoritism based on moral beliefs, they may be two sides of the same coin. Non-cooperative behavior violates social norms and is therefore closely related to moral behavior, which has been shown to elicit stronger in-group favoritism (Bilancini et al., 2020; Parker & Janoff-Bulman, 2013). Regardless, the preferential treatment of in-group individuals and cooperators violates the principle of impartiality inherent in utilitarian thinking (Kahane et al., 2018). Yet, they could still reflect less strict utilitarian cost-benefit calculations (Conway et al., 2018) that generate additional utility at the level of one's group.

Despite pronounced preferences for inaction (Baron & Ritov, 1994; Samuelson & Zeckhauser, 1988; Schaich Borg et al., 2006), trade-offs between participants surprisingly persist despite an option for random allocation and in contexts of withdrawal (see also Wilkinson et al., 2020). This suggests that decision makers reveal actual preferences. A potential alternative explanation is that people do not state preferences but simply distaste random allocation (Keren & Teigen, 2010) or perceive it as action rather than inaction (DeScioli et al., 2011). Yet, the ventilator is also most likely withdrawn from the disfavored patient among features that also elicit relatively high probabilities of withholding. Moreover, contrary to a preference to avoid responsibility for harming others (Leonhardt et al., 2011; Waldmann et al., 2012), the higher probability that favored patients receive the ventilator is not due to an increase in random allocations on disfavored patients. Thus, participants appear to not disguise their preferences by using random allocations (Batson et al., 1999; Dana et al., 2007; Dwenger et al., 2012) in triage decisions. Of course, that participants decided on behalf of another person (a doctor) could already have satisfied their goal of avoiding responsibility. Together this indicates that, at least in impersonal sacrificial dilemmas, decision makers weigh the benefits of allocating the ventilator to either individual with the cost of being responsible for harming others.

Differences in allocation based on patient features are in line with moral reasoning strategies that integrate information and are sensitive to cost-benefit analyses. This idea is in line with findings that small variations in the task (Christensen & Gomila, 2012) or context (Carnes et al., 2015; Decety & Cacioppo, 2012; FeldmanHall et al., 2012) can affect the responses to moral dilemmas. In-group effects indicate that information is integrated differently by different individuals. Moreover, Jin et al. (2021) show, those who are affected by the crisis and consider infections more seriously also weigh information differently. In fact, those who are more personally concerned might weigh patients' utility less strongly and aim to save the ones who are worst off.

Nevertheless, sensitivity to various kinds of information does not necessarily mean that the underlying process truly is a calculation of costs and benefits. Our data are still in line with various integration processes and may also involve simple rules (Gigerenzer, 2010) (e.g., to save the more healthy, more cooperative, more similar individual). Rules or processes used could also vary by participant and their perceptions of the crisis. In sum, our results underscore that various features of the decision situation affect moral decisions and hopefully spark future research which addresses the underlying processes more closely.

How people make moral decisions during a crisis that makes precisely these decisions more realistic may provide a valuable window into what affects moral decisions and reasoning. Beyond that, it has also been repeatedly argued that those who are affected should be asked about the ethical guidelines that are employed in making sacrificial decisions in the COVID-19 pandemic (Jin et al., 2021; Wilkinson et al., 2020) but also in the context of self-driving cars (Awad et al., 2018). However, it has also been argued that people's intuitions about moral issues should not enter ethical guidelines (Singer, 2005) and that asking the public will yield biases driven by majority views (Jaques, 2019). Although we find that people's decisions strongly reflect features that can be thought to reflect utilitarian principles and ethical guidelines, people also endorse features that are not covered by those guidelines. Our results suggest to side with those who caution against giving the public's moral judgments on hypothetical questions too much weight in informing ethical guidelines. Although these intuitions should enter the discourse to enhance the communication of the principles used (Fallucchi et al., 2020), it may be a more sustainable goal to invest in more fully understanding the reasoning process behind moral decisions in general.

# A

# Supplements for Chapter 2: Coordination under competition

## A.1 Example task

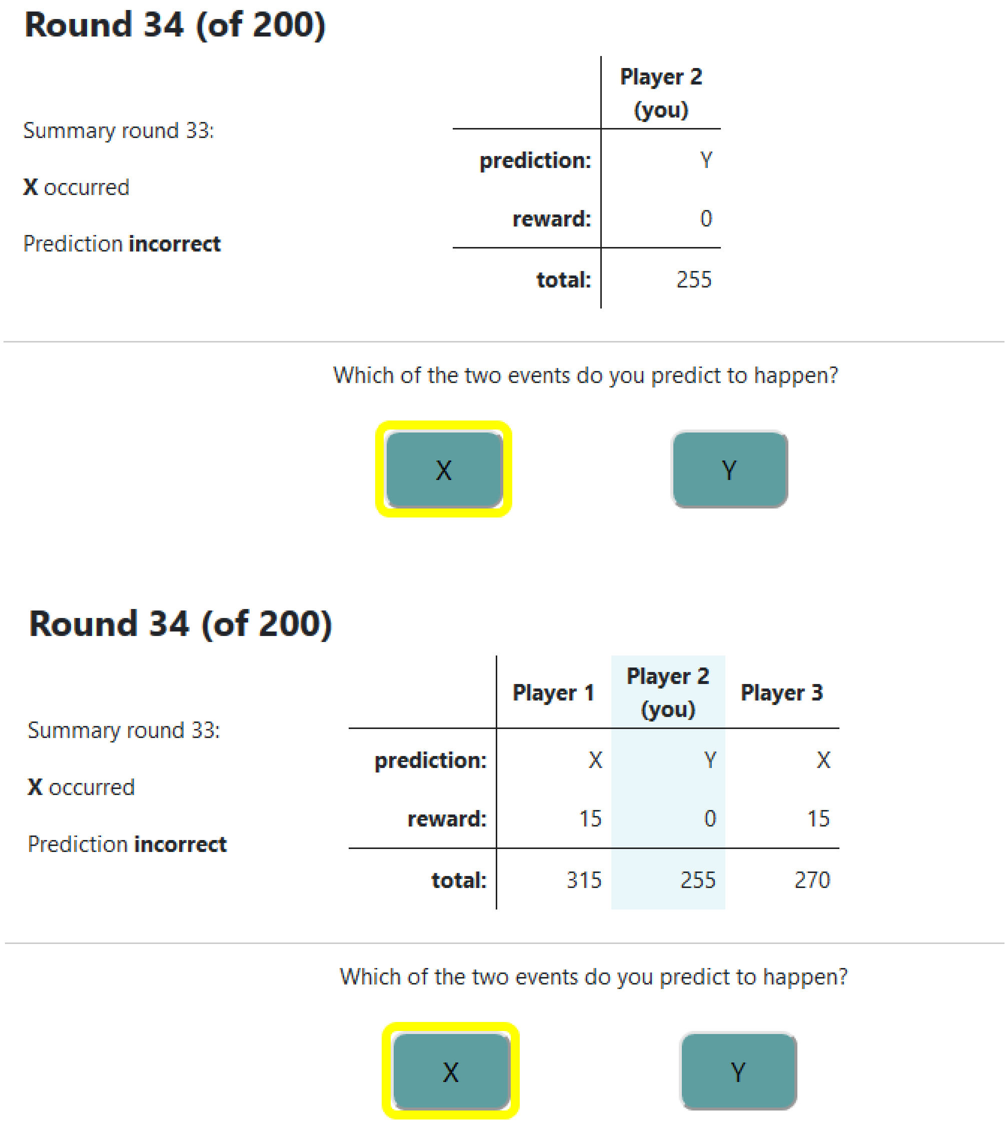

Round 34 (of 200)

Summary round 33:

**X** occurred

Prediction **incorrect**

| | Player 2 (you) |
|---|---|
| **prediction:** | Y |
| **reward:** | 0 |
| **total:** | 255 |

Which of the two events do you predict to happen?

Round 34 (of 200)

Summary round 33:

**X** occurred

Prediction **incorrect**

| | Player 1 | Player 2 (you) | Player 3 |
|---|---|---|---|
| **prediction:** | X | Y | X |
| **reward:** | 15 | 0 | 15 |
| **total:** | 315 | 255 | 270 |

Which of the two events do you predict to happen?

Figure A.1: Example scenarios as seen by the participants in the aware condition (upper panel) and observation condition (lower panel). The only difference between the display in the aware and unaware condition (not shown) was that in the unaware condition player number and "(you)" were not displayed.

## A.2 Description of the computerized agents

The computerized agents in Experiment 1 worked according to the Projective Simulation model (Briegel & De las Cuevas, 2012). The PS model in its simplest form involves an environment with percepts $s$ and actions $a$. In each discrete time-step $t$ each of $N$ agents $i$ (with $i \in 1, \ldots, N$) uses the current percept, accesses her memory of percept action pairs for this particular percept, and chooses the previously most rewarded action given the percept with some probability. These probabilities are derived from a random walk through memory, which is realized by the *hopping matrix* $h$, in which rows correspond to percepts and columns correspond to actions. Every time a percept-action pair is rewarded, the corresponding value in $h$ ($h$-value) is increased by the reward value $R$, increasing the probability of that particular action given the percept. Additionally, the $h$-matrix is dampened on each time step, according to a parameter $\gamma$, simulating the forgetting of information.

The updating of $h$-values after each time-step takes place according to the following process (Ried et al., 2019):

$$h_{ij}^{t+1} - 1 = (1-\gamma) * (h_{ij}^{t} - 1) + \begin{cases} R(t) & \text{for used transitions} \\ 0 & \text{for unused transitions} \end{cases}$$

In the version we used there is one stable percept corresponding to the the 2 available options and two possible actions $s_0$ and $s_1$. Rewards depend on whether the selected option is rewarded on the given round or not. Under competition, the reward of the actions depends on the number of other agents having chosen the same option, since rewards are shared.

Our agents were characterized by 2 main parameters: the greed or inverse temperature $\phi$ determining the degree of non-randomness in a softmax rule to transform $h$-values into action probabilities, and dampening or forgetting $\gamma$, according to which the learned values are decreased towards equality on each time-step.

The translation of $h$-values to choice probabilities is given by:

$$p_{ij} = \frac{h_{ij}^{\phi}}{\sum h_{ij}^{\phi}}$$

The agents we used were determined based on their behavior in preliminary simulations and both agents used $\phi = 0.35$. The agent type that learned relatively stable behavior without competition (maximizing), was defined by lower values of forgetting $\gamma = 0.05$, whereas the agent who produced behavior around probability matching was defined by higher values of forgetting $\gamma = 0.15$.

# A.3 Additional tables Experiment 1

## A.3.1 Additional descriptive statistics

Table A.1: Coordination under competition; Experiment 1: Cognitive Reflection Test (CRT) items by Mata (2016) and response distributions.

| | | Answer | | Percentage | | |
|---|---|---|---|---|---|---|
| Nr. | Item wording | correct | intuitive | correct | wrong intuitive | wrong non-intuitive |
| 1 | A TV and a DVD together cost 88 dollars. The TV costs 80 dollars more than the DVD. How much does the DVD cost? | 4 dollars | 8 dollars | 51.29 | 37.74 | 10.97 |
| 2 | If it takes 10 hens 10 days to lay 10 eggs, how long would it take 100 hens to lay 100 eggs? | 10 days | 100 days | 74.84 | 21.94 | 3.23 |
| 3 | A computer virus is spreading through the system of a computer. Every minute, the number of infected files doubles. If it takes 100 minutes for the virus to infect all of the system, how long would it take for the virus to infect half of the system? | 99 minutes | 50 minutes | 65.48 | 23.55 | 10.97 |

*Note:* $N = 310$.

Table A.2: Coordination under competition; Experiment 1: Descriptive statistics of the variables used in regression analysis and their correlations.

| | *M* | *SD* | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
|---|---|---|---|---|---|---|---|---|---|
| 1 Gender: Female[a] | 0.43 | 0.50 | — | **0.18** | -0.13 | -0.12 | 0.04 | 0.01 | -0.01 |
| 2 Age (years) | 36.85 | 11.24 | 0.00 | — | -0.02 | 0.07 | -0.00 | 0.02 | 0.02 |
| 3 Education: Bachelor or higher[a] | 0.53 | 0.50 | 0.02 | 0.75 | — | 0.05 | -0.02 | -0.00 | 0.02 |
| 4 CRT (items correct) | 1.92 | 1.15 | 0.03 | 0.19 | 0.38 | — | -0.02 | -0.08 | -0.05 |
| 5 Aware[a] | 0.32 | 0.47 | 0.50 | 1.00 | 0.77 | 0.70 | — | -0.46 | -0.02 |
| 6 Social[a] | 0.31 | 0.46 | 0.90 | 0.77 | 0.98 | 0.14 | 0.00 | — | 0.06 |
| 7 Competition[a] | 0.54 | 0.50 | 0.86 | 0.68 | 0.71 | 0.43 | 0.71 | 0.26 | 1.00 |

$N = 310$.

[a] Dummy-coded; CRT: Cognitive Reflection Test.

Pearson correlation coefficients are presented above the diagonal, and unadjusted p-values are presented below the diagonal. Only the correlation between age and gender is significant after Holm-adjusting the $p$-values and is displayed in bold.

### A.3.2 Model tables

Table A.3: Models predicting the probability that the human player selects the more likely option as a function of (1) round, (2) round and individual main effects and interactions of competition and social information, (3) adding the cross-level interactions of competition and social information with round, (4) splitting individuals in the competitive condition into those who experienced competition with bots and those that did not, and (5) adding interactions with player's CRT. The main text reports models (3), (4), and (5), although model (3) does not fit systematically better than model (1), using backwards elimination. All models control for non-independence of observations by adding random intercepts and time-slopes for individual human players (correlations omitted for sake of brevity).

| | *Dependent variable:* | | | | |
|---|---|---|---|---|---|
| | p(more likely option) | | | | |
| | (1) | (2) | (3) | (4) | (5) |
| Round$^{a}$ | 0.442$^{***}$ | 0.441$^{***}$ | 0.446$^{***}$ | 0.542$^{***}$ | 0.445$^{***}$ |
| | (0.041) | (0.041) | (0.040) | (0.059) | (0.040) |
| Competitive rewards$^{b}$ | | −0.034 | −0.186$^{**}$ | | −0.170$^{**}$ |
| | | (0.053) | (0.082) | | (0.080) |
| Aware$^{b}$ (Reference: no information) | | −0.041 | −0.142 | −0.220 | −0.129 |
| | | (0.075) | (0.116) | (0.165) | (0.114) |
| Observation$^{b}$ (Reference: no information) | | −0.074 | −0.086 | −0.041 | −0.069 |
| | | (0.076) | (0.118) | (0.173) | (0.116) |
| Competitive rewards×Aware | | −0.045 | 0.080 | | 0.096 |
| | | (0.075) | (0.116) | | (0.114) |
| Competitive rewards×Observation | | 0.071 | −0.044 | | −0.064 |
| | | (0.076) | (0.118) | | (0.117) |
| Round×Competitive rewards | | | −0.096$^{**}$ | | −0.096$^{**}$ |
| | | | (0.040) | | (0.040) |
| Round×Aware | | | −0.064 | −0.143$^{*}$ | −0.063 |
| | | | (0.056) | (0.082) | (0.056) |
| Round×Observation | | | −0.008 | 0.065 | −0.008 |
| | | | (0.057) | (0.086) | (0.057) |
| Round× Competitive rewards×Aware | | | 0.079 | | 0.079 |
| | | | (0.056) | | (0.056) |
| Round× Competitive rewards×Observation | | | −0.072 | | −0.073 |
| | | | (0.057) | | (0.057) |
| Competition (both bots maximize)$^{b}$ | | | | −0.887$^{***}$ | |
| | | | | (0.210) | |
| Competition (1 bot maximizes)$^{b}$ | | | | −0.034 | |
| | | | | (0.183) | |
| Round×Competition (both maximize) | | | | −0.286$^{***}$ | |
| | | | | (0.103) | |
| Round×Competition (1 maximizes) | | | | −0.129 | |
| | | | | (0.091) | |
| Aware×Competition (both maximize) | | | | 0.204 | |
| | | | | (0.294) | |
| Observation×Competition (both maximize) | | | | −0.025 | |
| | | | | (0.291) | |
| Aware×Competition (1 maximizes) | | | | 0.172 | |
| | | | | (0.261) | |
| Observation×Competition (1 maximizes) | | | | −0.029 | |
| | | | | (0.267) | |
| Round× Aware×Competition (both maximize) | | | | 0.136 | |
| | | | | (0.145) | |
| Round× Observation×Competition (both maximize) | | | | −0.129 | |
| | | | | (0.143) | |
| Round× Aware×Competition (1 maximizes) | | | | 0.182 | |
| | | | | (0.129) | |
| Round× Observation×Competition (1 maximizes) | | | | −0.140 | |
| | | | | (0.132) | |
| CRT score$^{c}$ | | | | | 0.152$^{***}$ |
| | | | | | (0.045) |
| Competitive rewards×CRT score | | | | | 0.011 |
| | | | | | (0.045) |
| Aware×CRT score | | | | | −0.002 |
| | | | | | (0.064) |
| Observation×CRT score | | | | | −0.014 |
| | | | | | (0.062) |
| Competitive rewards× Aware×CRT score | | | | | 0.016 |
| | | | | | (0.065) |
| Competitive rewards× Observation×CRT score | | | | | 0.083 |
| | | | | | (0.063) |
| Education (more than Bachelor)$^{d}$ | | | | | −0.023 |
| | | | | | (0.105) |
| Gender female$^{d}$ | | | | | 0.072 |
| | | | | | (0.107) |
| Age$^{c}$ | | | | | 0.008$^{*}$ |
| | | | | | (0.005) |
| Constant | 1.512$^{***}$ | 1.505$^{***}$ | 1.513$^{***}$ | 1.695$^{***}$ | 1.492$^{***}$ |
| | (0.084) | (0.083) | (0.082) | (0.118) | (0.112) |
| Number of Individuals | 310 | 310 | 310 | 310 | 310 |
| Random intercept | 1.432 | 1.421 | 1.4 | 1.362 | 1.363 |
| Random slopes round | 0.676 | 0.675 | 0.659 | 0.654 | 0.657 |
| Correlation intercept, slope | 0.794 | 0.793 | 0.786 | 0.788 | 0.788 |
| Observations | 62,000 | 62,000 | 62,000 | 62,000 | 62,000 |
| Log Likelihood | −30,737.940 | −30,736.040 | −30,731.030 | −30,721.050 | −30,721.720 |
| AIC | 61,485.880 | 61,492.090 | 61,492.070 | 61,484.110 | 61,491.450 |

*Note:* $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

$^{a}$: centered at 100, scaled to 50, $^{b}$: effect coded, $^{c}$: mean centered, $^{d}$: dummy coded

## A.4 Additional tables Experiment 2

### A.4.1 Additional descriptive statistics

Table A.4: Coordination under competition; Experiment 2: Cognitive Reflection Test (CRT) items by Mata (2016) and response distributions.

| | | Answer | | Percentage | | |
|---|---|---|---|---|---|---|
| Nr. | Item wording | correct | intuitive | correct | wrong intuitive | wrong non-intuitive |
| 1 | A TV and a DVD together cost 88 dollars. The TV costs 80 dollars more than the DVD. How much does the DVD cost? | 4 dollars | 8 dollars | 45.95 | 48.22 | 5.83 |
| 2 | If it takes 10 hens 10 days to lay 10 eggs, how long would it take 100 hens to lay 100 eggs? | 10 days | 100 days | 68.28 | 28.16 | 3.56 |
| 3 | A computer virus is spreading through the system of a computer. Every minute, the number of infected files doubles. If it takes 100 minutes for the virus to infect all of the system, how long would it take for the virus to infect half of the system? | 99 minutes | 50 minutes | 60.84 | 28.32 | 10.84 |

*Note:* $N = 618$.

Table A.5: Coordination under competition; Experiment 2: Descriptive statistics of the variables used in regression analysis and their correlations.

| | *M* | *SD* | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
|---|---|---|---|---|---|---|---|---|---|
| 1 Gender: Female[a] | 0.47 | 0.50 | — | **0.13** | 0.01 | **-0.21** | -0.07 | 0.03 | -0.01 |
| 2 Age (years) | 36.40 | 11.28 | 0.00 | — | 0.04 | 0.05 | 0.02 | -0.07 | 0.09 |
| 3 Education: Bachelor or higher[a] | 0.57 | 0.50 | 0.73 | 0.32 | — | 0.11 | 0.03 | -0.04 | 0.05 |
| 4 CRT (items correct) | 1.75 | 1.15 | 0.00 | 0.22 | 0.01 | — | 0.05 | -0.02 | 0.02 |
| 5 Aware[a] | 0.28 | 0.45 | 0.10 | 0.64 | 0.45 | 0.18 | — | -0.48 | -0.05 |
| 6 Social[a] | 0.37 | 0.48 | 0.48 | 0.10 | 0.36 | 0.60 | 0.00 | — | -0.03 |
| 7 Competition[a] | 0.50 | 0.50 | 0.82 | 0.03 | 0.21 | 0.64 | 0.22 | 0.49 | 1.00 |

$N = 618$.

[a] Dummy-coded; CRT: Cognitive Reflection Test.

Pearson correlation coefficients are presented above the diagonal, and unadjusted p-values are presented below the diagonal. Only the correlations of gender with age and CRT are significant after Holm-adjusting the *p*-values and is displayed in bold.

## A.4.2 Model tables

Table A.6: Models of the probability that a group achieves a 2:1 distribution (IFD) in a given round as a function of (1) round, (2) adding individual main effects of competition and social information, (3) adding the interaction of competition and social information, (4) adding the cross-level interactions of competition and social information with round, and (5) adding interactions with the lowest CRT score in the group. Thereby controlling for non-independence of observations by adding random intercepts and time-slopes for groups.

| | *Dependent variable:* | | | | |
|---|---|---|---|---|---|
| | p(IFD reached) | | | | |
| | (1) | (2) | (3) | (4) | (5) |
| Round$^{a}$ | −0.031 | −0.030 | −0.030 | −0.029 | −0.019 |
| | (0.023) | (0.023) | (0.023) | (0.022) | (0.029) |
| Aware$^{b}$ (Reference: No information) | | −0.010 | −0.017 | 0.013 | −0.004 |
| | | (0.049) | (0.048) | (0.055) | (0.068) |
| Observation$^{b}$ (Reference: No information) | | 0.028 | 0.031 | 0.022 | 0.031 |
| | | (0.046) | (0.045) | (0.051) | (0.066) |
| Competitive rewards$^{b}$ | | 0.131$^{***}$ | 0.122$^{***}$ | 0.200$^{***}$ | 0.133$^{***}$ |
| | | (0.035) | (0.034) | (0.037) | (0.047) |
| Aware×Competitive rewards | | | −0.104$^{**}$ | −0.107$^{*}$ | −0.209$^{***}$ |
| | | | (0.048) | (0.055) | (0.068) |
| Observation×Competitive rewards | | | 0.113$^{**}$ | 0.118$^{**}$ | 0.070 |
| | | | (0.045) | (0.051) | (0.066) |
| Round×Aware | | | | 0.036 | 0.023 |
| | | | | (0.032) | (0.042) |
| Round×Observation | | | | −0.012 | 0.004 |
| | | | | (0.030) | (0.040) |
| Round×Competitive rewards | | | | 0.095$^{***}$ | 0.112$^{***}$ |
| | | | | (0.022) | (0.029) |
| Round×<br>Aware×Competitive rewards | | | | −0.004<br>(0.032) | −0.002<br>(0.042) |
| Round×<br>Observation×Competitive rewards | | | | 0.006<br>(0.030) | −0.029<br>(0.040) |
| Lowest CRT (in group)$^{c}$ | | | | | −0.072$^{*}$ |
| | | | | | (0.040) |
| Round×Lowest CRT | | | | | −0.018 |
| | | | | | (0.025) |
| Aware×Lowest CRT | | | | | 0.063 |
| | | | | | (0.056) |
| Observation×Lowest CRT | | | | | −0.057 |
| | | | | | (0.058) |
| Competitive rewards×Lowest CRT | | | | | 0.090$^{**}$ |
| | | | | | (0.040) |
| Round×<br>Aware×Lowest CRT | | | | | 0.020<br>(0.034) |
| Round×<br>Observation×Lowest CRT | | | | | −0.028<br>(0.036) |
| Round×<br>Competitive rewards×Lowest CRT | | | | | −0.020<br>(0.025) |
| Aware×<br>Competitive rewards×Lowest CRT | | | | | 0.107$^{*}$<br>(0.056) |
| Observation×<br>Competitive rewards×Lowest CRT | | | | | 0.104$^{*}$<br>(0.058) |
| Round×Aware×<br>Competitive rewards×Lowest CRT | | | | | −0.003<br>(0.034) |
| Round×Observation×<br>Competitive rewards×Lowest CRT | | | | | 0.056<br>(0.036) |
| Constant | −0.580$^{***}$ | −0.582$^{***}$ | −0.583$^{***}$ | −0.582$^{***}$ | −0.534$^{***}$ |
| | (0.040) | (0.038) | (0.038) | (0.037) | (0.047) |
| Number of Groups | 206 | 206 | 206 | 206 | 206 |
| Random intercept | 0.572 | 0.538 | 0.53 | 0.524 | 0.495 |
| Random time slopes | 0.063 | 0.315 | 0.315 | 0.3 | 0.296 |
| Correlation intercept, slope | 0.553 | 0.514 | 0.519 | 0.503 | 0.511 |
| Observations | 123,600 | 123,600 | 123,600 | 123,600 | 123,600 |
| Log Likelihood | −77,343.040 | −77,335.960 | −77,332.380 | −77,322.760 | −77,307.760 |
| AIC | 154,696.100 | 154,687.900 | 154,684.800 | 154,675.500 | 154,669.500 |

*Note:* $^{*}$p<0.1; $^{**}$p<0.05; $^{***}$p<0.01

$^{a}$: centered at 100, scaled to 50, $^{b}$: effect coded, $^{c}$: mean centered, $^{d}$: dummy coded

Table A.7: Models predicting the probability that a player selects the more likely option as a function of (1) round, (2) adding individual main effects of competition and social information, (3) adding the interaction of competition and social information, and (4) adding the cross-level interactions of competition and social information with round. Thereby controlling for non-independence of observations by adding random intercepts and time-slopes for participants nested in groups.

| | *Dependent variable:* | | | | |
|---|---|---|---|---|---|
| | p(more likely option) | | | | |
| | (1) | (2) | (3) | (4) | (5) |
| Round$^{a}$ | 0.575$^{***}$<br>(0.036) | 0.573$^{***}$<br>(0.036) | 0.573$^{***}$<br>(0.036) | 0.576$^{***}$<br>(0.036) | 0.592$^{***}$<br>(0.036) |
| Aware$^{b}$ (Reference: No information) | | 0.018<br>(0.080) | 0.020<br>(0.079) | 0.077<br>(0.108) | 0.064<br>(0.108) |
| Observation$^{b}$ (Reference: No information) | | −0.026<br>(0.074) | −0.029<br>(0.074) | −0.096<br>(0.101) | −0.116<br>(0.102) |
| Competitive rewards$^{b}$ | | −0.083<br>(0.055) | −0.078<br>(0.055) | −0.259$^{***}$<br>(0.073) | −0.254$^{***}$<br>(0.074) |
| Aware×Competitive rewards | | | 0.065<br>(0.079) | 0.113<br>(0.108) | 0.115<br>(0.108) |
| Observation×Competitive rewards | | | −0.022<br>(0.074) | −0.108<br>(0.101) | −0.121<br>(0.102) |
| Round×Aware | | | | 0.041<br>(0.051) | 0.035<br>(0.052) |
| Round×Observation | | | | −0.047<br>(0.048) | −0.048<br>(0.049) |
| Round×Competitive rewards | | | | −0.129$^{***}$<br>(0.035) | −0.128$^{***}$<br>(0.036) |
| Round× Aware×Competitive rewards | | | | 0.035<br>(0.051) | 0.032<br>(0.052) |
| Round× Observation×Competitive rewards | | | | −0.062<br>(0.048) | −0.081$^{*}$<br>(0.049) |
| CRT score$^{c}$ | | | | | 0.255$^{***}$<br>(0.062) |
| Round×CRT score | | | | | 0.070$^{**}$<br>(0.029) |
| Aware×CRT score | | | | | 0.073<br>(0.091) |
| Observation×CRT score | | | | | −0.112<br>(0.085) |
| Competitive rewards×CRT score | | | | | −0.013<br>(0.062) |
| Round×Aware×CRT score | | | | | 0.017<br>(0.043) |
| Round×Observation×CRT score | | | | | −0.026<br>(0.041) |
| Round× Competitive rewards×CRT score | | | | | −0.004<br>(0.029) |
| Aware× Competitive rewards×CRT score | | | | | −0.114<br>(0.091) |
| Observation× Competitive rewards×CRT score | | | | | 0.042<br>(0.085) |
| Round×Aware× Competitive rewards×CRT score | | | | | −0.037<br>(0.043) |
| Round×Observation× Competitive rewards×CRT score | | | | | −0.037<br>(0.041) |
| Constant | 1.741$^{***}$<br>(0.075) | 1.742$^{***}$<br>(0.074) | 1.744$^{***}$<br>(0.074) | 1.748$^{***}$<br>(0.074) | 1.801$^{***}$<br>(0.075) |
| Number of Groups | 206 | 206 | 206 | 206 | 206 |
| Random intercept (participant) | 1.747 | 1.734 | 1.733 | 1.711 | 1.681 |
| Random round slopes (participant) | 0.16 | 0.794 | 0.793 | 0.781 | 0.773 |
| Correlation intercept, slope (Participant) | 0.82 | 0.817 | 0.816 | 0.811 | 0.81 |
| Random intercept (group) | 0.243 | 0.218 | 0.215 | 0.251 | 0.243 |
| Random round slopes (group) | 0.029 | 0.159 | 0.161 | 0.138 | 0.143 |
| Correlation intercept, slope (Group) | -1 | -1 | -1 | -1 | -1 |
| Observations | 123,307 | 123,307 | 123,307 | 123,307 | 123,307 |
| Log Likelihood | −57,736.490 | −57,735.330 | −57,734.990 | −57,726.420 | −57,713.180 |
| AIC | 115,489.000 | 115,492.700 | 115,496.000 | 115,488.800 | 115,486.400 |

*Note:* $^{*}$p<0.1; $^{**}$p<0.05; $^{***}$p<0.01

$^{a}$: centered at 100, scaled to 50, $^{b}$: effect coded, $^{c}$: mean centered, $^{d}$: dummy coded

# B

# Supplements for Chapter 3: Recommendation and exploration

## B.1 Methods and results Simulations

We first simulated individuals from the function learning model in Wu et al. (2018). Simulated individuals estimate a function over the value of available options and attach an estimate of uncertainty to each. They integrate value and uncertainty by preferentially sampling options with high value and high uncertainty (optimistic sampling, Auer et al., 2002; Reverdy et al., 2014). The function learning model and the choice function involve different parameters and weights.

$$\mathrm{UCB}(\mathbf{x}) = m_t(\mathbf{x}) + \beta s_t(\mathbf{x})$$

According to the UCB strategy, an option's expected value $m_t(\mathbf{x})$ and an option's uncertainty $s_t(\mathbf{x})$ (from the Gaussian Process function, not shown), are integrated to an option's utility. Unchosen options have larger uncertainty and are therefore more likely chosen if an individual has a non-zero value for the exploration bonus $\beta$, resulting in directed exploration. The recommended option $\mathbf{x_a}$ receives a fixed bonus $|\mu(\mathbf{x_a})|$ on this expected utility from the trial the recommendation is accepted.

Based on this estimate, decision probabilities of each option were determined by a softmax rule, in which a higher temperature parameter $\tau$ is related to more random exploration:

$$p(x) = \frac{\exp(UCB(x)/\tau)}{\sum_{j=1}^{121} \exp(UCB(j)/\tau)}$$

If models were localized, the option utilities $\mathrm{UCB}(\mathbf{x})$ were weighted by the inverse of the Manhattan Distance (IMD) to the last revealed tile $\mathrm{IMD}(x, x_{prev}) = \sum_{n=1}^{121} = (|x_i - x_{prev}|)^{-1}$. For $x_1 = x_{prev}$, IMD was set to 1.

To conduct our simulations, we sampled individuals' parameters from the parameter estimates from the study by Wu and colleagues. Simulated individuals explored the spatially correlated bandits and received a recommendation after 5 or 20 trials ("clicks"). Then, they either continued exploration or exploit the recommendation. We compared, how following or exploiting the recommendation instead of continuing exploration affected (a) the overall points accumulated and (b) the probability of finding the best option (global maximum) in the environment.

The simulation findings are summarized in Figure 3.3 (see main text). If individuals' goal is to accumulate rewards, they can benefit from following the recommendation (Panel A). Individuals' average rewards are never lower than without a recommendation (colored lines). This comes at the cost of a decreased chance of finding the best option (Panel B). Accordingly, a simple heuristic that always takes and exploits a recommendation can already be expected to be fairly robust with respect to accumulating rewards and also finding the global maximum for short time horizons.

However, individuals can, but do not need to, rely on simple heuristics. They may also use generalization to inform their decision to take advice and to integrate the recommendation with their previous experience. Introducing a utility bonus to function learning, yields a model that gives recommendations a priority but still integrates them with previous information, reflecting the potentially rewarding nature of following recommendations (Biele et al., 2009).

$$\mathrm{UCB}_{bonus}(\mathbf{x}) = m_t(\mathbf{x}) + \beta s_t(\mathbf{x}) + |\mu(\mathbf{x_a})|$$

We simulated a total of 10,000 individuals. We varied individuals' learning parameters and whether their choices were affected by a penalty for more distant options. This penalty leads them to more likely stick close to already encountered options (localization). Simulated individuals acquired previous experience of 5 or 20 clicks. Then, starting from this situation, they received one randomly determined, unrevealed tile as recommendation. Recommendations were either familiar or novel. This was determined based on the uncertainty $s_t(\mathbf{x})$ assigned to the options by a basic function learning model (see Figure 3.1C and D). For each individual, behavior after this recommendation was simulated under bonus values ranging from 0 (ignoring the recommendation) to 0.5 (high recommendation weight).

## B.2 Providing recommendations: Model table

Table B.1: Average MD to the previous option as a function of whether individuals are asked to provide a recommendation after exploration. Model 3 excludes the maximization condition and involves 2 contrasts comparing (a) providing a recommendation (with or without rewards depending on recipient performance) against just accumulating points and (b) providing a recommendation and being rewarded according to performance against providing a recommendation and being rewarded according to accumulated points. Correlations between random intercepts are not displayed.

| | *Dependent variable:* | | |
|---|---|---|---|
| | average MD to previous (exploration locality) | | |
| | (1) | (2) | (3) |
| C1: All others vs. Accumulation | 0.073$^{***}$ | 0.074$^{***}$ | |
| | (0.016) | (0.016) | |
| C2: Maximizing vs. Recommendation (both) | −0.093$^{***}$ | −0.094$^{***}$ | |
| | (0.023) | (0.023) | |
| C3: Recommendation only vs. recommendation plus accumulation | 0.180$^{***}$ | 0.184$^{***}$ | |
| | (0.040) | (0.040) | |
| SVO (effect coded, reference: prosocial) | 0.035 | 0.035 | |
| | (0.076) | (0.076) | |
| Contrast recommendation vs. accumulation | | | 0.065$^{***}$ |
| | | | (0.023) |
| Contrast recommendation only vs. accumulation | | | 0.183$^{***}$ |
| | | | (0.040) |
| Environment type (effect coded, reference: smooth) | 0.219$^{***}$ | 0.219$^{***}$ | 0.221$^{***}$ |
| | (0.060) | (0.060) | (0.067) |
| Risk preference (z-standardized) | −0.019 | −0.019 | −0.016 |
| | (0.072) | (0.072) | (0.073) |
| Round (centered at 4) | 0.012 | 0.012 | −0.004 |
| | (0.012) | (0.012) | (0.015) |
| Order condition Maximizing first (dummy coded, reference: Accumulation first) | −0.210 | −0.211 | |
| | (0.219) | (0.219) | |
| Order condition Accumulation plus recommendation first (dummy coded, reference: Accumulation first) | −0.110 | −0.110 | |
| | (0.201) | (0.201) | |
| Order condition Recommendation only first (dummy coded, reference: Accumulation first) | 0.156 | 0.156 | |
| | (0.207) | (0.207) | |
| C1×SVO | | −0.006 | |
| | | (0.016) | |
| C2×SVO | | 0.008 | |
| | | (0.023) | |
| C3×SVO | | −0.020 | |
| | | (0.040) | |
| Constant | 3.092$^{***}$ | 3.093$^{***}$ | 2.983$^{***}$ |
| | (0.161) | (0.161) | (0.093) |
| Number of Individuals | 154 | 154 | 154 |
| Intercept for individual | 0.817 | 0.822 | 0.822 |
| Intercept for environment | 0.257 | 0.234 | 0.235 |
| Residual | 0.973 | 0.973 | 0.974 |
| Observations | 1,232 | 1,232 | 924 |
| Log Likelihood | −1,890.477 | −1,898.570 | −1,436.970 |
| AIC | 3,808.955 | 3,831.140 | 2,891.941 |
| *Note:* | | $^{*}$p<0.1; $^{**}$p<0.05; $^{***}$p<0.01 | |

## B.3 Receiving recommendations: Model selection and model tables

### B.3.1 Clicking on the recommendation

Table B.2: Models predicting the probability of clicking on the recommendation immediately after receiving it with (1) main effects and (2) interactions and models predicting clicking the recommendation anytime during the round with (3) main effects and (4) interactions. Correlations between random intercepts are not displayed.

| | *Dependent variable:* | | | |
|---|---|---|---|---|
| | p(clicked immediately) | | p(clicked anytime) | |
| | (1) | (2) | (3) | (4) |
| Option value$^{a}$ | $-0.014$ | $-0.013$ | $0.043^{**}$ | $0.146^{***}$ |
| | (0.014) | (0.014) | (0.020) | (0.054) |
| Familiarity$^{b}$ | $0.164^{***}$ | $0.157^{***}$ | $0.185^{***}$ | $0.333^{*}$ |
| (reference: novel) | (0.048) | (0.049) | (0.068) | (0.173) |
| Recipient experience 15 clicks$^{b}$ | $-0.003$ | $-0.004$ | $-0.050$ | $-0.056$ |
| (reference: 5 clicks) | (0.048) | (0.049) | (0.068) | (0.077) |
| Environment type$^{b}$ (reference: smooth) | $0.142^{**}$ | $0.133^{**}$ | $0.205^{***}$ | $0.474^{**}$ |
| | (0.059) | (0.061) | (0.071) | (0.199) |
| Option value×Familiarity | | 0.017 | | $0.050^{*}$ |
| | | (0.014) | | (0.027) |
| Option value×Recipient experience | | $0.039^{***}$ | | 0.025 |
| | | (0.014) | | (0.023) |
| Familiarity×Recipient experience | | $-0.108^{**}$ | | $-0.261^{***}$ |
| | | (0.049) | | (0.079) |
| Option value×Environment type | | $-0.007$ | | $-0.032$ |
| | | (0.015) | | (0.028) |
| Familiarity×Environment type | | $-0.106^{**}$ | | $-0.011$ |
| | | (0.051) | | (0.093) |
| Recipient experience× | | 0.044 | | 0.008 |
| Environment type | | (0.051) | | (0.081) |
| Option value×Familiarity× | | 0.006 | | $-0.021$ |
| Recipient experience | | (0.014) | | (0.023) |
| Option value×Familiarity× | | $-0.021$ | | $-0.076^{***}$ |
| Environment type | | (0.015) | | (0.025) |
| Option value×Recipient experience× | | $0.034^{**}$ | | $0.057^{**}$ |
| Environment type | | (0.015) | | (0.025) |
| Familiarity×Recipient experience× | | 0.020 | | 0.118 |
| Environment type | | (0.052) | | (0.082) |
| Option value×Familiarity× | | $-0.021$ | | $-0.006$ |
| Recipient experience×Environment type | | (0.015) | | (0.024) |
| Constant | $1.109^{***}$ | $1.121^{***}$ | $3.044^{***}$ | $3.637^{***}$ |
| | (0.115) | (0.118) | (0.185) | (0.310) |
| Number of Individuals | 353 | 353 | 353 | 353 |
| Intercept for individual | 1.792 | 1.816 | 2.068 | 2.466 |
| Intercept for grid | 0.135 | 0.152 | – | – |
| Value slopes for individual | – | – | – | 0.19 |
| Familiarity slopes for individual | – | – | – | 0.365 |
| Environment slopes for individual | – | – | – | 0.625 |
| Observations | 2,824 | 2,824 | 2,824 | 2,824 |
| Log Likelihood | $-1{,}497.433$ | $-1{,}484.176$ | $-917.407$ | $-892.079$ |
| AIC | 3,008.866 | 3,004.352 | 1,846.814 | 1,836.159 |

*Note:* $^{*}$p<0.1; $^{**}$p<0.05; $^{***}$p<0.01

$^{a}$: Centered at 0.65 and scaled by 0.15; $^{b}$: effect coded, $^{c}$: dummy coded

**Immediate taking** The fit of a logistic generalized mixed model of the probability of clicking on the recommendation in a given grid directly after receiving it (immediate taking) is improved by the inclusion of random intercepts for individuals, $\chi^2(1) = 547.481$, $p < 0.001$, $\Delta AIC = 545.48$, and environments $\chi^2(1) = 4.031$, $p = 0.045$, $\Delta AIC = 2.03$. This indicates that the probability of clicking on the recommendation varies by individual and environment.

Model fit is further improved by inclusion of round-level (level of observation) fixed effects for the factors of our design (recommendation high vs. low, recommendation familiarity, recipient experience, environment type) and their interactions, $\chi^2(15) = 43.409$, $p < 0.001$, $\Delta AIC = 13.41$. Inclusion of random slopes for all fixed effects does not improve model fit further on the individual level, $\chi^2(14) = 4.871$, $p = 0.988$, $\Delta AIC = -23.13$, or the environment level, $\chi^2(14) = 7.936$, $p = 0.893$, $\Delta AIC = -20.06$.

**Cumulative taking** The fit of a logistic generalized mixed model of the probability of clicking on the recommendation in a given grid anytime after receiving it (cumulative taking) is improved by the inclusion of random intercepts for individuals, $\chi^2(2) = 389.601$, $p < 0.001$, $\Delta AIC = 385.60$, but not for environments, $\chi^2(1) = 2.405$, $p = 0.121$, $\Delta AIC = 0.41$. This indicates that the probability of clicking on the recommendation varies by individual. Therefore, we include random intercepts for individuals.

Model fit is further improved by inclusion of round-level (level of observation) fixed effects for the factors of our design (recommendation high vs. low, recommendation familiarity, recipient experience, and environment type), as well as their interactions, $\chi^2(16) = 48.956$, $p < 0.001$, $\Delta AIC = 16.96$. Inclusion of individual-level random slopes for fixed effects of recommendation quality, familiarity, and environment type (smooth vs. rough) additionally improves model fit, $\chi^2(14) = 4.871$, $p = 0.988$, $\Delta AIC = -23.13$. Therefore, we will interpret a model with these specifications.

### B.3.2 Global maximum

The fit of a logistic generalized mixed model of the probability of finding the best option in the environment in a given grid is improved by the inclusion of random intercepts for individuals, $\chi^2(1) = 79.677$, $p < 0.001$, $\Delta AIC = 77.68$, and environments, $\chi^2(2) = 267.329$, $p < 0.001$, $\Delta AIC = 263.33$. This indicates that the probability of finding the best option varies by individual and environment.

Model fit is further improved by inclusion of round-level (level of observation) fixed effects for the factors of our design (recommendation value, option familiarity, dummy coded indicators for rounds without recommendations and genuine recommendations, recipient experience), environment type and their interactions, $\chi^2(23) = 295.314$, $p < 0.001$, $\Delta AIC = 249.31$. Also the interaction of the indicators for rounds with no recommendation and genuine recommendations with the value of that recommendation (or its no recommendation control) improve model fit, $\chi^2(8) = 41.953$, $p < 0.001$, $\Delta AIC = 25.95$. Inclusion of random slopes for all fixed effects does not improve model fit further on the individual level, $\chi^2(19) = 0.000$, $p = 1.000$, $\Delta AIC = -64.42$, or the environment level, $\chi^2(19) = 0.000$, $p = 1.000$, $\Delta AIC = -61.61$.

### B.3.3 Distance to recommendation

The fit of a linear mixed model of the MD to the recommendation (or its control option) after receiving (finding) it is improved by the inclusion of random intercepts for individuals, $\chi^2(1) = 70.674$, $p < 0.001$, $\Delta AIC = 68.67$ and environments $\chi^2(1) = 15.452$,

Table B.3: Models predicting predicting the probability that the largest option in the grid is found as a function of (1) recommendation conditions and (2) their interactions. Correlations between random intercepts are not displayed.

| | *Dependent variable:* | |
|---|---|---|
| | p(best option found) | |
| | (1) | (2) |
| Option value$^{a}$ | $0.105^{***}$ | $0.071^{***}$ |
| | (0.009) | (0.013) |
| Familiarity$^{b}$ (reference: novel) | $-0.057$ | $-0.054$ |
| | (0.042) | (0.042) |
| No advice round$^{c}$ | $-0.112$ | $-0.179^{*}$ |
| | (0.095) | (0.102) |
| Genuine advice$^{c}$ | $0.247^{**}$ | $-0.025$ |
| | (0.105) | (0.128) |
| Recipient experience | $0.370^{***}$ | $0.393^{***}$ |
| | (0.035) | (0.043) |
| Environment type | 0.145 | 0.140 |
| | (0.121) | (0.125) |
| Option value×Familiarity | | $-0.026^{**}$ |
| | | (0.013) |
| Option value×Recipient experience | | 0.015 |
| | | (0.013) |
| Familiarity×Recipient experience | | $-0.038$ |
| | | (0.042) |
| Option value×Environment type | | $0.039^{***}$ |
| | | (0.013) |
| Familiarity×Environment type | | $-0.046$ |
| | | (0.043) |
| Recipient experience×Environment type | | $-0.086^{**}$ |
| | | (0.043) |
| Option value×No advice | | 0.044 |
| | | (0.028) |
| Recipient experience×No advice | | $-0.025$ |
| | | (0.102) |
| Environment type×No advice | | $-0.048$ |
| | | (0.104) |
| Option value×Genuine advice | | $0.116^{***}$ |
| | | (0.025) |
| Recipient experience×Genuine advice | | $-0.146$ |
| | | (0.120) |
| Environment type×Genuine advice | | 0.081 |
| | | (0.130) |
| Option value×Familiarity×Recipient experience | | $-0.016$ |
| | | (0.013) |
| Option value×Familiarity×Environment type | | 0.005 |
| | | (0.013) |
| Option value×Recipient experience×Environment type | | $-0.016$ |
| | | (0.013) |
| Familiarity×Recipient experience×Environment type | | 0.004 |
| | | (0.043) |
| Option value×Recipient experience×No advice | | $-0.084^{***}$ |
| | | (0.028) |
| Option value×Environment type×No advice | | $-0.080^{***}$ |
| | | (0.028) |
| Recipient experience×Environment type×No advice | | 0.023 |
| | | (0.104) |
| Option value×Recipient experience×Genuine advice | | $-0.003$ |
| | | (0.025) |
| Option value×Environment type×Genuine advice | | $-0.031$ |
| | | (0.025) |
| Recipient experience×Environment type×Genuine advice | | $0.228^{*}$ |
| | | (0.122) |
| Option value×Familiarity×Recipient experience×Environment type | | $-0.004$ |
| | | (0.013) |
| Option value×Recipient experience×Environment type×No advice | | 0.029 |
| | | (0.028) |
| Option value×Recipient experience×Environment type×Genuine advice | | $-0.018$ |
| | | (0.025) |
| Constant | $-0.435^{***}$ | $-0.423^{***}$ |
| | (0.128) | (0.130) |
| Number of Individuals | 353 | 353 |
| Intercept for individual | 0.676 | 0.698 |
| Intercept for grid | 0.516 | 0.523 |
| Observations | 4,236 | 4,236 |
| Log Likelihood | −2,626.582 | −2,592.651 |
| AIC | 5,271.164 | 5,253.302 |

*Note:* $^{*}$p<0.1; $^{**}$p<0.05; $^{***}$p<0.01

$^{a}$: centered (mean: 0.65, scale: 0.15), $^{b}$: effect coded, 0 for no-advice and genuine conditions, $^{c}$: dummy coded

$p < 0.001$, $\Delta AIC = 13.45$. This indicates that the distance to the previously selected option varies by individual and environment.

Model fit is further improved by inclusion of round-level (level of observation) fixed effects for the factors of our design (recommendation value, option familiarity, dummy coded indicators for rounds without recommendations and genuine recommendations, recipient experience), environment type and their interactions, $\chi^2(7) = 1058.634$, $p < 0.001$, $\Delta AIC = 1044.63$. The interaction of the indicators for rounds with no recommendation and genuine recommendations with the value of that recommendation (or its no recommendation control) additionally improves model fit, $\chi^2(24) = 69.751$, $p < 0.001$, $\Delta AIC = 21.75$. Inclusion of individual-level random slopes for all fixed effects improves model fit, $\chi^2(27) = 101.390$, $p < 0.001$, $\Delta AIC = 47.39$, whereas inclusion on the environment level does not, $\chi^2(27) = 11.391$, $p = 0.996$, $\Delta AIC = -42.61$.

### B.3.4 Distance to previous selection

The fit of a linear mixed model of the MD to the previously selected option after a recommendation was received in a given grid is improved by the inclusion of random intercepts for individuals, $\chi^2(1) = 1252.577$, $p < 0.001$, $\Delta AIC = 1250.58$ and environments $\chi^2(1) = 35.087$, $p < 0.001$, $\Delta AIC = 33.09$. This indicates that the distance to the previously selected option varies by individual and environment.

Model fit is further improved by inclusion of round-level (level of observation) fixed effects for the factors of our design (recommendation value, option familiarity, dummy coded indicators for rounds without recommendations and genuine recommendations, recipient experience), environment type and their interactions, $\chi^2(23) = 288.319$, $p < 0.001$, $\Delta AIC = 242.32$. The interaction of the indicators for rounds with no recommendation and genuine recommendations with the value of that recommendation (or its no recommendation control) does not further improve model fit, $\chi^2(8) = 6.118$, $p = 0.634$, $\Delta AIC = -9.88$, Inclusion of individual-level random slopes for all fixed effects improves model fit, $\chi^2(27) = 56.613$, $p = 0.001$, $\Delta AIC = 2.61$, whereas inclusion on the environment level does not, $\chi^2(27) = 9.034$, $p = 1.000$, $\Delta AIC = -44.97$.

Table B.4: Models predicting the average MD to the recommendation from (1) main effects of recommendation conditions and (2) their interactions, as well as models predicting the MD to the previous click from (3) main effects of recommendation conditions and (4) their interactions. Correlations between random intercepts are not displayed.

| | *Dependent variable:* | | | |
|---|---|---|---|---|
| | MD to recommendation | | MD to recommendation | |
| | (1) | (2) | (3) | (4) |
| Option value[a] | $-0.337^{***}$ | $-0.367^{***}$ | $-0.069^{***}$ | $-0.068^{***}$ |
| | (0.011) | (0.016) | (0.005) | (0.005) |
| Familiarity[b] | $-0.462^{***}$ | $-0.467^{***}$ | $-0.053^{***}$ | $-0.053^{***}$ |
| | (0.048) | (0.051) | (0.019) | (0.019) |
| No-advice condition[c] | $-0.836^{***}$ | $-0.789^{***}$ | 0.017 | 0.013 |
| | (0.107) | (0.111) | (0.045) | (0.046) |
| Genuine condition[c] | 0.075 | −0.148 | 0.059 | 0.061 |
| | (0.116) | (0.128) | (0.048) | (0.049) |
| Recipient experience[d] | $0.094^{**}$ | $0.118^{***}$ | −0.013 | 0.011 |
| 15 trials (reference: 5 trials) | (0.039) | (0.046) | (0.018) | (0.021) |
| Rough environment | $-0.233^{***}$ | $-0.268^{***}$ | $0.084^{***}$ | $0.097^{***}$ |
| (reference: smooth) | (0.058) | (0.064) | (0.026) | (0.029) |
| Option value×Familiarity | | $0.028^{**}$ | | 0.008 |

| | | | | |
|---|---|---|---|---|
| | | (0.013) | | (0.006) |
| Option value×Recipient experience | | 0.012 | | 0.006 |
| | | (0.013) | | (0.004) |
| Familiarity×Recipient experience | | 0.081* | | −0.039** |
| | | (0.045) | | (0.019) |
| Option value×Rough environment | | 0.037*** | | −0.006 |
| | | (0.014) | | (0.004) |
| Familiarity×Rough environment | | 0.213*** | | 0.057*** |
| | | (0.048) | | (0.020) |
| Experience 15 trials× | | −0.016 | | −0.021 |
| Rough environment | | (0.048) | | (0.020) |
| Option value×No-advice | | 0.040 | | |
| | | (0.028) | | |
| Experience 15 trials×No-advice | | −0.204** | | −0.107** |
| | | (0.104) | | (0.042) |
| Rough environment×No-advice | | 0.065 | | −0.048 |
| | | (0.106) | | (0.045) |
| Option value×Genuine | | 0.092*** | | |
| | | (0.025) | | |
| Experience 15 trials×Genuine | | −0.108 | | −0.049 |
| | | (0.120) | | (0.045) |
| Rough environment×Genuine | | 0.255** | | −0.013 |
| | | (0.130) | | (0.050) |
| Option value×Familiarity× | | 0.004 | | 0.001 |
| Experience 15 trials | | (0.013) | | (0.006) |
| Option value×Familiarity× | | 0.010 | | −0.001 |
| Rough environment | | (0.014) | | (0.006) |
| Option value×Experience 15 trials× | | 0.004 | | 0.001 |
| Rough environment | | −0.014 | | 0.001 |
| | | (0.014) | | (0.004) |
| Familiarity×Experience 15 trials× | | 0.004 | | 0.001 |
| Rough environment | | −0.040 | | −0.010 |
| | | (0.048) | | (0.020) |
| Option value×Experience 15 trials× | | −0.009 | | |
| No-advice | | (0.027) | | |
| Option value×Rough environment× | | −0.059** | | |
| No-advice | | (0.028) | | |
| Experience 15 trials×Rough environment× | | −0.038 | | −0.013 |
| No-advice | | (0.107) | | (0.046) |
| Option value×Experience 15 trials× | | 0.033 | | |
| Genuine | | (0.024) | | |
| Option value×Rough environment× | | −0.055** | | |
| Genuine | | (0.025) | | |
| Experience 15 trials×Rough environment× | | 0.093 | | −0.0002 |
| Genuine | | (0.125) | | (0.047) |
| Option value×Familiarity× | | 0.025* | | 0.005 |
| Experience 15 trials×Rough environment | | (0.014) | | (0.006) |
| Option value×Experience 15 trials× | | 0.002 | | |
| Rough environment×No-advice | | (0.028) | | |
| Option value×Experience 15 trials× | | 0.008 | | |
| Rough environment×Genuine | | (0.025) | | |
| Constant | 4.783*** | 4.798*** | 2.165*** | 2.166*** |
| | (0.078) | (0.080) | (0.053) | (0.053) |
| Number of Individuals | 353 | 353 | 353 | 353 |
| Random intercept individual | 0.838 | 0.947 | 0.85 | 0.851 |
| Random slopes option value (individual) | 0.189 | 0.152 | 0.039 | 0.039 |
| Random slopes familiarity (individual) | – | 0.467 | 0.047 | 0.055 |
| Random slopes experience (individual) | – | 0.146 | 0.158 | 0.161 |
| Random slopes environment (individual) | – | 0.141 | 0.088 | 0.089 |
| Random slopes no-advice (individual) | – | 0.744 | 0.303 | 0.322 |
| Random slopes genuine advice (individual) | – | 0.497 | 0.208 | 0.208 |
| Random intercept environment | 0.189 | 0.19 | 0.091 | 0.092 |
| Residual variance | 2.538 | 2.373 | 0.999 | 0.993 |
| Observations | 4,236 | 4,236 | 4,236 | 4,236 |
| Log Likelihood | −10,110.450 | −10,025.750 | −6,530.143 | −6,514.861 |
| AIC | 20,240.910 | 20,175.500 | 13,134.290 | 13,137.720 |

*Notes:* $^{*}$p<0.1; $^{**}$p<0.05; $^{***}$p<0.01

$^{a}$: centered (mean: 0.65, scale: 0.15), $^{b}$: effect coded, 0 for no-advice and genuine conditions, $^{c}$: dummy coded

Random effect correlations not shown.

# C

# Supplements for Chapter 4: Allocation of scarce resources

## C.1 Patient features and their rationale

Table C.1: Patient features and potential reasons to prefer the second individual.

| Feature | Disfavored | Favored | Utility increase through |
|---|---|---|---|
| *Related to the greater good* | | | |
| age | 68 years old | 42 years old | years of life saved, risk factor related to survival probability |
| profession health | police officer | health professional | instrumental value for saving lives |
| children: 32-year-old woman | no children | young children | benefits to mother and children |
| physical fitness | excess-weight | healthy-weight | risk factor related to survival probability |
| alcohol consumption | excess drinking | moderate drinking | risk factor related to survival probability |
| hereditary diabetes | hereditary diabetes | otherwise healthy | risk factor related to survival probability |
| *Not obviously related to the greater good* | | | |
| sex | male | female | |
| ethnicity | black | white | saving in-group individual (among blacks/whites) |
| living in the U.S | recently arrived in the U.S. with a work permit | born and raised in the U.S. | saving in-group individual |
| residence permit | recently took refuge in the U.S | born and raised in the U.S. | saving in-group individual |
| profession police | business executive | police officer | instrumental value |
| religious beliefs | does not believe in God | believes in God | saving in-group individual among the religious |
| children: 64-year-old woman | without children | adult children | contribution to society, dependent others |
| relationship status | single | married | contribution to society, dependent others |
| disability | healthy person in a wheelchair | otherwise healthy person | Life-years adjusted for quality (QALYs) |
| financial honesty | evades paying taxes | regularly pays due taxes | cooperative behavior and altruistic punishment |
| charity-related behavior | did not donate to charity last year | donated to charity last year | cooperative behavior and altruistic punishment |
| organ donor status | refuses to donate organs | is willing to donate organs | cooperative behavior and altruistic punishment |
| vaccination behavior | did not get the flu shot this season | got the flu shot this season | cooperative behavior and altruistic punishment |

## C.2 Example task

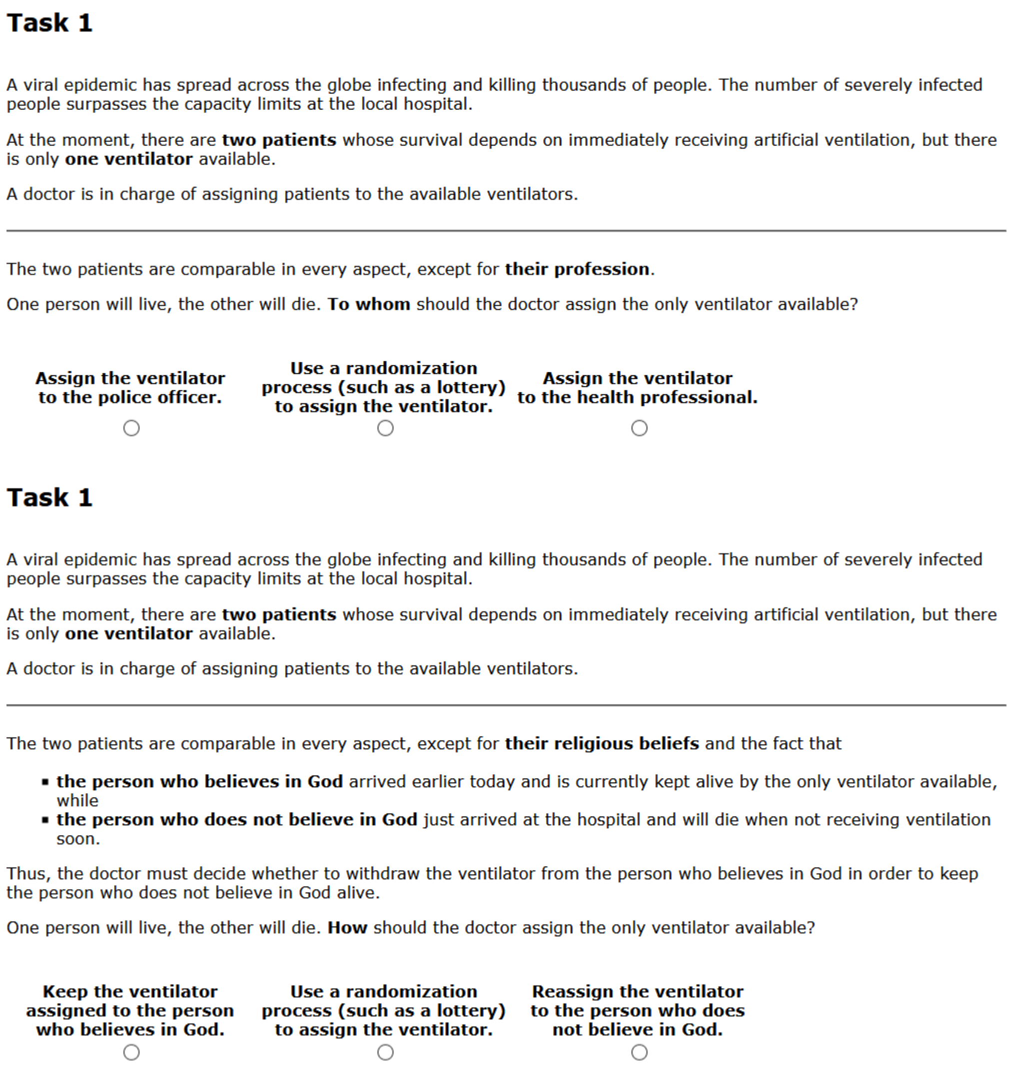

**Task 1**

A viral epidemic has spread across the globe infecting and killing thousands of people. The number of severely infected people surpasses the capacity limits at the local hospital.

At the moment, there are **two patients** whose survival depends on immediately receiving artificial ventilation, but there is only **one ventilator** available.

A doctor is in charge of assigning patients to the available ventilators.

The two patients are comparable in every aspect, except for **their profession**.

One person will live, the other will die. **To whom** should the doctor assign the only ventilator available?

**Assign the ventilator to the police officer.** ○

**Use a randomization process (such as a lottery) to assign the ventilator.** ○

**Assign the ventilator to the health professional.** ○

**Task 1**

A viral epidemic has spread across the globe infecting and killing thousands of people. The number of severely infected people surpasses the capacity limits at the local hospital.

At the moment, there are **two patients** whose survival depends on immediately receiving artificial ventilation, but there is only **one ventilator** available.

A doctor is in charge of assigning patients to the available ventilators.

The two patients are comparable in every aspect, except for **their religious beliefs** and the fact that

- **the person who believes in God** arrived earlier today and is currently kept alive by the only ventilator available, while
- **the person who does not believe in God** just arrived at the hospital and will die when not receiving ventilation soon.

Thus, the doctor must decide whether to withdraw the ventilator from the person who believes in God in order to keep the person who does not believe in God alive.

One person will live, the other will die. **How** should the doctor assign the only ventilator available?

**Keep the ventilator assigned to the person who believes in God.** ○

**Use a randomization process (such as a lottery) to assign the ventilator.** ○

**Reassign the ventilator to the person who does not believe in God.** ○

Figure C.1: Example scenarios as seen by the participants in the baseline condition without a default (upper panel) and with a default (lower panel).

## C.3 Sample description

| | Wave | | | |
|---|---|---|---|---|
| | 1 (n = 225) | 2 (n = 225) | 3 (n = 659) | 4 (n = 644) |
| **Condition** | | | | |
| non-preferred default | 0 | 0 | 222 | 215 |
| no default | 225 | 225 | 213 | 213 |
| preferred default | 0 | 0 | 224 | 216 |
| **Demographics** | | | | |
| Female | 45.3 | 44.4 | 41.3 | 39.1 |
| Age | 37.7 | 38.5 | 37.2 | 38.1 |
| Religiosity | 316.0 | 341.8 | 359.0 | 340.5 |
| Political | 462.2 | 476.0 | 452.7 | 464.3 |
| **Any children** | | | | |
| Any children | 35.1 | 43.1 | 43.7 | 40.8 |
| **Education** | | | | |
| none | 0.4 | 0.4 | 0.2 | 0.2 |
| high school | 25.3 | 20.9 | 16.9 | 18.5 |
| post high | 13.8 | 12.4 | 11.4 | 11.3 |
| vocational | 2.7 | 3.6 | 3.0 | 2.5 |
| BA | 48.0 | 44.0 | 50.8 | 49.7 |
| MA | 9.3 | 15.6 | 15.4 | 16.0 |
| PhD | 0.4 | 3.1 | 2.3 | 1.9 |
| **Income (in US-$)** | | | | |
| $\leq$ 5000 | 3.1 | 7.1 | 3.8 | 5.9 |
| $<$ 5000 | 5.8 | 4.4 | 4.4 | 4.3 |
| $<$ 10000 | 13.8 | 8.9 | 13.2 | 13.4 |
| $<$ 15000 | 9.8 | 8.4 | 7.3 | 9.2 |
| $<$ 25000 | 16.9 | 12.9 | 13.4 | 12.7 |
| $<$ 35000 | 23.6 | 16 | 19.3 | 18.8 |
| $<$ 50000 | 20.4 | 34.7 | 30.7 | 26.1 |
| $<$ 80000 | 6.7 | 7.6 | 8 | 9.6 |
| **Ethnicity** | | | | |
| White/Caucasian | 78.2 | 80.9 | 71.1 | 76.9 |
| African American | 8.9 | 9.3 | 14.1 | 10.6 |
| Asian American | 7.1 | 4.4 | 7.3 | 6.8 |
| American Indian or Alaskan Native | 0.4 | 0.9 | 0.3 | 0.5 |
| Native Hawaiian or Pacific Islander | 0.0 | 0.0 | 0.0 | 0.2 |
| Latino | 3.6 | 2.7 | 5.2 | 3.7 |
| Other | 1.8 | 1.8 | 2.0 | 1.4 |
| **Marital status** | | | | |
| single | 44.0 | 38.7 | 40.7 | 37.9 |
| partner | 16.0 | 14.2 | 9.9 | 12.7 |
| married | 40.0 | 46.7 | 48.8 | 48.1 |
| not say | 0.0 | 0.4 | 0.6 | 1.2 |
| **Time in the U.S.** | | | | |
| less than 20 years | 6.2 | 7.6 | 10.0 | 7.0 |
| migrated recently | 4.9 | 4.5 | 5.6 | 6.7 |
| born in the U.S. | 88.9 | 87.9 | 84.4 | 86.3 |
| **Residence area (population)** | | | | |
| village (¡ 3000) | 4.9 | 6.2 | 5.0 | 4.3 |
| small town (3000 to 15,000) | 13.3 | 11.6 | 15.2 | 14.9 |
| town (15,000 to 100,000) | 28.4 | 32.9 | 32.2 | 34.3 |
| city (100,000 to 1,000,000) | 35.6 | 32.9 | 32.2 | 27.2 |
| large city (¿ 1,000,000) | 17.8 | 15.6 | 15.0 | 18.6 |
| don't know | 0.0 | 0.9 | 0.5 | 0.6 |
| **Standard drinks** | | | | |
| none | 47.1 | 36.4 | 34.0 | 38.4 |
| 1-2 | 32.4 | 41.3 | 39.3 | 37.1 |
| 3-4 | 13.3 | 17.8 | 17.6 | 15.2 |
| 5-6 | 6.7 | 2.7 | 6.2 | 6.2 |
| 7-9 | 0.4 | 0.9 | 2.0 | 1.9 |
| 10+ | 0.0 | 0.9 | 0.9 | 1.2 |
| **BMI** | | | | |
| 0 to 18.5 | 4.9 | 7.1 | 7.7 | 6.5 |
| 18.5 to 25 | 40.0 | 46.7 | 47.2 | 45.3 |
| 25 to 30 | 32.9 | 29.3 | 30.3 | 30.9 |
| $>$ 30 | 22.2 | 16.9 | 14.7 | 17.2 |
| **Own health** | | | | |
| excellent | 27.1 | 27.6 | 26.0 | 21.3 |
| good | 51.1 | 56.9 | 60.6 | 64.2 |
| fair | 20.9 | 13.8 | 12.2 | 12.8 |
| poor | 0.9 | 1.8 | 1.2 | 1.7 |
| **Own COVID infection** | | | | |
| can't tell | 8.4 | 4.4 | 5.3 | 7.5 |
| no | 88.0 | 91.6 | 88.8 | 87.4 |
| not confirmed | 1.8 | 1.8 | 2.9 | 3.1 |
| yes | 1.8 | 2.2 | 3.0 | 2.0 |
| **Someone with COVID known** | | | | |
| can't tell | 4.4 | 4.4 | 4.2 | 5.0 |
| no | 86.2 | 80.4 | 76.8 | 76.6 |
| not confirmed | 3.6 | 4.9 | 7.3 | 6.1 |
| yes | 5.8 | 10.2 | 11.7 | 12.4 |
| **Decision preferences** | | | | |
| Larger later | 28.4 | 34.2 | 37.0 | 33.7 |
| Risky option | 10.7 | 13.3 | 15.2 | 13.5 |
| Altruism (amount) | 133.3 | 166.6 | 179.1 | 178.8 |

## C.4 Model tables

Table C.3: Modelling survival of the favored patient as a function of the experimentally varied factors (1) feature type and (2) feature, as well as default condition and collection wave.

| | *Dependent variable:* | |
|---|---|---|
| | logit(survival) | |
| | (1) | (2) |
| Favored default (Expected to be favored has ventilator) | 1.073$^{***}$ | 1.006$^{***}$ |
| | (0.051) | (0.086) |
| Disfavored default (Expected to be disfavored has ventilator) | −1.169$^{***}$ | −1.290$^{***}$ |
| | (0.050) | (0.085) |
| Data collection wave 2 | −0.036 | −0.036 |
| (Reference wave 1) | (0.071) | (0.071) |
| Data collection wave 3 | −0.058 | −0.058 |
| (Reference wave 1) | (0.054) | (0.054) |
| Data collection wave 4 | 0.030 | 0.030 |
| (Reference wave 1) | (0.054) | (0.054) |
| Unrelated feature (Feature unrelated to saving more individuals) | −0.226$^{***}$ | |
| | (0.068) | |
| Unrelated feature×Favored default | 0.080$^{***}$ | |
| | (0.017) | |
| Unrelated feature×Disfavored default | −0.069$^{***}$ | |
| | (0.016) | |
| *Single feature effects: survival probability of favored* | | |
| *(second entry compared to average of all features):* | | |
| Sex: | | −0.483$^{***}$ |
| Male vs. female | | (0.049) |
| Ethnicity: | | −0.545$^{***}$ |
| Black vs. white | | (0.049) |
| U.S. residence: | | 0.113$^{**}$ |
| Recent arrival with work permit vs. born and raised | | (0.051) |
| Residence permit status: | | 0.208$^{***}$ |
| Recent refugee vs. person born in the U.S. | | (0.052) |
| Profession health: | | 0.233$^{***}$ |
| Police officer vs. health professional | | (0.052) |
| Profession police: | | −0.095$^{*}$ |
| Business executive vs. police officer | | (0.050) |
| Religious beliefs: | | −0.604$^{***}$ |
| Unreligious vs. religious | | (0.049) |
| Young mother: | | 0.751$^{***}$ |
| 32-year old woman without vs. with young children | | (0.057) |
| Children: | | 0.170$^{***}$ |
| 64-year-old woman without vs. with children | | (0.051) |
| Relationship status: | | −0.077 |
| Not married vs. married | | (0.050) |
| Physical fitness | | 0.173$^{***}$ |
| Excess-weight vs. healthy-weight person | | (0.051) |
| Health wheelchair | | −0.377$^{***}$ |
| Person in a wheelchair vs. person not in a wheelchair | | (0.049) |
| Alcohol consumption | | 0.477$^{***}$ |
| Excess drinking vs. moderate drinking | | (0.054) |
| Health diabetes | | −0.061 |
| Person with diabetes vs. person without | | (0.050) |
| Financial honesty | | 0.305$^{***}$ |
| Evades paying taxes vs. pays due taxes | | (0.052) |
| Charity-related behavior | | −0.190$^{***}$ |
| Did not donate vs. donated last year | | (0.050) |
| Organ donation | | −0.029 |
| Not registered vs. registered | | (0.050) |
| Flu vaccination | | −0.351$^{***}$ |
| No flu shot vs. flu shot | | (0.049) |
| *Single feature interactions with default condition:* | | |
| Sex×Favored default | | 0.330$^{***}$ |
| | | (0.096) |
| Ethnicity×Favored default | | 0.143 |
| | | (0.094) |
| U.S. residence×Favored default | | 0.081 |
| | | (0.101) |
| Residence permit status×Favored default | | −0.142 |
| | | (0.100) |
| Profession health×Favored default | | −0.007 |
| | | (0.102) |
| Profession police×Favored default | | 0.024 |
| | | (0.097) |
| Religious beliefs×Favored default | | 0.038 |

| | | |
|---|---|---|
| | | (0.093) |
| Young mother×Favored default | | −0.191* |
| | | (0.110) |
| Children×Favored default | | −0.042 |
| | | (0.100) |
| Relationship status×Favored default | | 0.089 |
| | | (0.098) |
| Physical fitness×Favored default | | −0.146 |
| | | (0.099) |
| Health×Favored default | | 0.038 |
| | | (0.095) |
| Alcohol consumption×Favored default | | −0.116 |
| | | (0.105) |
| Health diabetes×Favored default | | −0.099 |
| | | (0.097) |
| Financial honesty×Favored default | | −0.038 |
| | | (0.102) |
| Charity-related behavior×Favored default | | 0.111 |
| | | (0.097) |
| Organ donation×Favored default | | −0.027 |
| | | (0.098) |
| Flu vaccination×Favored default | | 0.126 |
| | | (0.096) |
| Sex×Disfavored default | | −0.062 |
| | | (0.095) |
| Ethnicity×Disfavored default | | −0.172* |
| | | (0.096) |
| U.S. residence×Disfavored default | | 0.122 |
| | | (0.094) |
| Residence permit status×Disfavored default | | 0.205** |
| | | (0.094) |
| Profession health×Disfavored default | | −0.064 |
| | | (0.094) |
| Profession police×Disfavored default | | 0.020 |
| | | (0.093) |
| Religious beliefs×Disfavored default | | −0.026 |
| | | (0.095) |
| Young mother×Disfavored default | | −0.097 |
| | | (0.098) |
| Children×Disfavored default | | −0.220** |
| | | (0.094) |
| Relationship status×Disfavored default | | −0.012 |
| | | (0.093) |
| Physical fitness×Disfavored default | | 0.015 |
| | | (0.094) |
| Health×Disfavored default | | −0.098 |
| | | (0.094) |
| Alcohol consumption×Disfavored default | | 0.036 |
| | | (0.096) |
| Health diabetes×Disfavored default | | 0.364*** |
| | | (0.093) |
| Financial honesty×Disfavored default | | 0.134 |
| | | (0.095) |
| Charity-related behavior×Disfavored default | | −0.210** |
| | | (0.094) |
| Organ donation×Disfavored default | | −0.020 |
| | | (0.093) |
| Flu vaccination×Disfavored default | | −0.145 |
| | | (0.095) |
| Constant | 0.611*** | 0.620*** |
| | (0.079) | (0.041) |
| Number of Individuals | 1753 | 1753 |
| Random intercepts for participants | 1.17 | 1.173 |
| Random intercepts for patient features | 0.27 | – |
| Observations | 33,301 | 33,301 |
| Log Likelihood | −28,410.830 | −28,332.880 |
| AIC | 56,843.670 | 56,781.770 |
| *Note:* | All categorical variables are effect-coded.<br>*p<0.1; **p<0.05; ***p<0.01 | |

Table C.4: Modelling survival of the favored patient in decision situations without default as a function of decision maker control variables, items, their interactions of interest, and collection wave. Patient features are effect-coded effects compare the item to the average of all patient features. Interactions between item indicators and decision maker features indicate in-group effects.

| | *Dependent variable:* | |
|---|---|---|
| | logit(survival) | |
| | (1) | (2) |
| Unrelated feature (Feature unrelated to saving more individuals) | $-0.176^{**}$ | $-0.240^{***}$ |
| | (0.073) | (0.058) |
| Data collection wave 2 (Reference wave 1) | $-0.034$ | $-0.029$ |
| | (0.034) | (0.034) |
| Data collection wave 3 (Reference wave 1) | 0.013 | 0.020 |
| | (0.035) | (0.035) |
| Data collection wave 4 (Reference wave 1) | $-0.041$ | $-0.041$ |
| | (0.035) | (0.035) |
| *Decision maker features:* | | |
| County-level COVID incidence$^{a}$ | 0.024 | |
| | (0.035) | |
| Perceived COVID outbreak severity$^{a}$ | $-0.002$ | |
| | (0.026) | |
| Perceived severity of COVID infection$^{a}$ | $-0.033$ | |
| | (0.024) | |
| Perceived probability of COVID infection$^{a}$ | $-0.028$ | |
| | (0.025) | |
| Self reported COVID infection: Can't tell (reference none) | 0.126 | |
| | (0.077) | |
| Self reported COVID infection: Not confirmed (reference none) | $-0.115$ | |
| | (0.101) | |
| Self reported COVID infection: Yes | $-0.205^{*}$ | |
| | (0.108) | |
| Decision maker sex | $-0.014$ | $-0.014$ |
| | (0.020) | (0.021) |
| Decision maker age (centered) | 0.0001 | $-0.0001$ |
| | (0.002) | (0.002) |
| Income larger than \$50k | 0.031 | 0.031 |
| | (0.020) | (0.021) |
| Ethnicity: Afro American (reference: Caucasian) | $-0.084^{*}$ | $-0.087^{*}$ |
| | (0.046) | (0.046) |
| Ethnicity: Other (reference: Caucasian) | $0.096^{**}$ | $0.110^{**}$ |
| | (0.044) | (0.045) |
| City or larger ($> 100,000$ inhabitants) | 0.002 | 0.003 |
| | (0.020) | (0.020) |
| Region: Midwest (reference: West) | 0.029 | 0.011 |
| | (0.041) | (0.041) |
| Region: Northeast (reference: West) | $-0.033$ | $-0.024$ |
| | (0.044) | (0.040) |
| Region: South (reference: West) | $0.059^{*}$ | $0.073^{**}$ |
| | (0.032) | (0.032) |
| Religiosity$^{a}$ | $-0.055^{**}$ | $0.084^{***}$ |
| | (0.022) | (0.032) |
| Liberal political orientation$^{a}$ | $-0.044^{**}$ | $-0.192^{***}$ |
| | (0.022) | (0.028) |
| Positive vaccination attitude$^{a}$ | | $0.100^{***}$ |
| | | (0.035) |
| Got flu shot (reference: did not get flu shot) | | $0.067^{*}$ |
| | | (0.037) |
| Altruism$^{b}$ | | $-0.033^{*}$ |
| | | (0.017) |
| BMI $< 18.5$ | | $-0.162$ |
| (underweight) | | (0.100) |
| BMI $25 - 30$ | | 0.046 |
| (overweight, reference: normal-weight) | | (0.056) |
| BMI $> 30$ | | $-0.015$ |
| (adipositas, reference: normal-weight) | | (0.067) |
| Registered organ donor (reference: non-registered) | | $0.097^{***}$ |
| | | (0.031) |
| Moderate drinking (reference: excess drinking) | | 0.059 |
| | | (0.038) |
| *Patient features (items):* | | |
| Residence: U.S. residence or residence permit status | | $0.173^{**}$ |
| | | (0.073) |

| | | |
|---|---|---|
| Religiosity | | −0.168$^{*}$ |
| | | (0.095) |
| Vaccination | | −0.074 |
| | | (0.095) |
| Charity-related behavior | | −0.004 |
| | | (0.095) |
| Financial honesty | | 0.224$^{**}$ |
| | | (0.095) |
| Organ donation | | 0.065 |
| | | (0.095) |
| Physical fitness | | −0.119 |
| | | (0.100) |
| Alcohol consumption | | 0.062 |
| | | (0.099) |
| Ethnicity (saving white person) | | −0.488$^{**}$ |
| | | (0.196) |
| *Patient feature × Decision maker feature-interactions:* | | |
| Residence×Liberal political orientation$^{a}$ | | −0.141$^{***}$ |
| | | (0.020) |
| Vaccination×Positive vaccination attitude$^{a}$ | | 0.058$^{**}$ |
| | | (0.029) |
| Vaccination×Got flu shot | | 0.068$^{**}$ |
| | | (0.031) |
| Charity-related behavior×Altruism$^{b}$ | | 0.017 |
| | | (0.011) |
| Financial honesty×Altruism$^{b}$ | | −0.030$^{***}$ |
| | | (0.011) |
| Organ donation×Registered organ donor | | 0.087$^{***}$ |
| | | (0.026) |
| Physical fitness×BMI $< 18.5$ | | −0.122 |
| | | (0.085) |
| Physical fitness×BMI $25 - 30$ | | −0.001 |
| | | (0.048) |
| Physical fitness×BMI $> 30$ | | −0.001 |
| | | (0.056) |
| Alcohol consumption (moderate)×Moderate drinking | | 0.022 |
| | | (0.033) |
| Ethnicity×Afro American | | −0.005 |
| | | (0.113) |
| Ethnicity×Other ethnicity | | −0.199$^{*}$ |
| | | (0.111) |
| Religiosity×Religiosity$^{a}$ | | 0.165$^{***}$ |
| | | (0.025) |
| Unrelated feature×COVID incidence | 0.015 | |
| | (0.020) | |
| Unrelated feature×Outbreak severity | −0.005 | |
| | (0.016) | |
| Unrelated feature×Infection severity | 0.049$^{***}$ | |
| | (0.015) | |
| Unrelated feature×Infection probability | −0.039$^{**}$ | |
| | (0.016) | |
| Unrelated feature×Own infection (can't tell) | −0.071 | |
| | (0.048) | |
| Unrelated feature×Own infection (not confirmed) | 0.066 | |
| | (0.063) | |
| Unrelated feature×Own infection (yes) | 0.064 | |
| | (0.065) | |
| Constant | 0.493$^{***}$ | 0.772$^{**}$ |
| | (0.086) | (0.313) |
| Number of Individuals | 876 | 876 |
| Random intercepts for participants | 0.459 | 0.463 |
| Random intercepts for patient features | 0.264 | 0.166 |
| Observations | 16,639 | 16,639 |
| Log Likelihood | −14,802.630 | −14,731.330 |
| AIC | 29,669.260 | 29,558.650 |

*Note:* $^{a}$: $z$-standardized, $^{b}$: mean-centered, all categorical variables are effect-coded
$^{*}$p<0.1; $^{**}$p<0.05; $^{***}$p<0.01

Table C.5: Modelling the probability of random allocations as a function of (1) demographics and indicators of actual and perceived of crisis severity and (2) demographics and their interactions with features of interest.

| | *Dependent variable:* | |
|---|---|---|
| | logit(random allocation) | |
| | (1) | (2) |
| Data collection wave 2 | 0.156 | 0.121 |
| (Reference wave 1) | (0.108) | (0.102) |
| Data collection wave 3 | $-0.271^{**}$ | $-0.140$ |
| (Reference wave 1) | (0.110) | (0.104) |
| Data collection wave 4 | 0.008 | $-0.007$ |
| (Reference wave 1) | (0.110) | (0.104) |
| Perceived severity of COVID infection$^{a}$ | | $-0.313^{***}$ |
| | | (0.085) |
| Perceived probability of COVID infection$^{a}$ | | 0.039 |
| | | (0.071) |
| Police officer | | $-0.235$ |
| Business executive vs. police-officer | | (0.319) |
| Health professional | | 0.385 |
| Police officer vs. health professional | | (0.336) |
| Self reported COVID infection: | | $0.718^{***}$ |
| Can't tell (reference none) | | (0.225) |
| Self reported COVID infection: | | $-0.523^{*}$ |
| Not confirmed (reference none) | | (0.299) |
| Self reported COVID infection: Yes | | $-0.288$ |
| | | (0.315) |
| Unrelated patient feature | | $0.823^{***}$ |
| (Feature unrelated to saving more individuals) | | (0.164) |
| Decision maker sex (female, reference: male) | | $0.245^{***}$ |
| | | (0.061) |
| Decision maker age$^{b}$ | | $0.017^{***}$ |
| | | (0.005) |
| Income larger than \$50k | | $-0.068$ |
| | | (0.061) |
| Ethnicity: Afro American (reference: Caucasian) | | $-0.108$ |
| | | (0.138) |
| Ethnicity: Other )reference: Caucasian) | | $-0.118$ |
| | | (0.131) |
| City or larger ($> 100,000$ inhabitants) | | 0.023 |
| | | (0.060) |
| Region: Midwest (reference: West) | | $-0.079$ |
| | | (0.120) |
| Region: Northeast (reference: West) | | $0.199^{*}$ |
| | | (0.120) |
| Region: South (reference: West) | | $-0.067$ |
| | | (0.095) |
| Religiosity$^{a}$ | | $-0.340^{***}$ |
| | | (0.067) |
| Liberal political orientation$^{a}$ | | $0.180^{***}$ |
| | | (0.065) |
| Police officer× | | $-0.083^{*}$ |
| Perceived severity of COVID infection | | (0.043) |
| Health professional× | | $-0.171^{***}$ |
| Perceived severity of COVID infection | | (0.045) |

| | | |
|---|---|---|
| Unrelated feature× | | −0.116*** |
| Perceived severity of COVID infection | | (0.025) |
| Unrelated feature× | | −0.033 |
| Perceived probability of COVID infection | | (0.025) |
| Constant | 0.311*** | −0.162 |
| | (0.063) | (0.463) |
| *Single feature effects: survival probability of favored* | | |
| *(second entry compared to average of all features):* | | |
| Sex: | 1.619*** | |
| Male vs. female | (0.097) | |
| Ethnicity: | 1.933*** | |
| Black vs. white | (0.104) | |
| U.S. residence: | −0.182** | |
| Recent arrival with work permit vs. born and raised | (0.080) | |
| Residence permit status: | −0.264*** | |
| Recent refugee vs. person born in the U.S. | (0.080) | |
| Profession health: | −0.361*** | |
| Police officer vs. health professional | (0.080) | |
| Profession police: | 0.052 | |
| Business executive vs. police officer | (0.081) | |
| Religious beliefs: | 1.260*** | |
| Unreligious vs. religious | (0.091) | |
| Young mother: | −1.310*** | |
| 32-year old woman without vs. with young children | (0.084) | |
| Older mother: | −0.101 | |
| 64-year-old woman without vs. with children | (0.080) | |
| Relationship status: | 0.455*** | |
| Not married vs. married | (0.083) | |
| Physical fitness: | −0.934*** | |
| Excess-weight vs. healthy-weight person | (0.081) | |
| Health (wheelchair): | 0.371*** | |
| Person in a wheelchair vs. person not in a wheelchair | (0.082) | |
| Alcohol consumption: | −0.998*** | |
| Excess drinking vs. moderate drinking | (0.082) | |
| Health (diabetes): | −0.939*** | |
| Person with diabetes vs. person without | (0.082) | |
| Financial honesty: | −0.546*** | |
| Evades paying taxes vs. pays due taxes | (0.080) | |
| Charity-related behavior: | 0.945*** | |
| Did not donate vs. donated last year | (0.087) | |
| Organ donation: | 0.073 | |
| Not registered vs. registered | (0.081) | |
| Vaccination: | 0.578*** | |
| No flu shot vs. flu shot | (0.084) | |
| Number of Individuals | 876 | 876 |
| Random intercepts for participant | 1.746 | 1.614 |
| Random intercepts for patient features | – | 0.608 |
| Observations | 16,639 | 16,639 |
| Log Likelihood | −8,674.088 | −8,643.647 |
| AIC | 17,394.180 | 17,345.290 |

*Note:* $^{a}$: $z$-standardized, $^{b}$: mean-centered, all categorical variables are effect-coded
$^{*}$p<0.1; $^{**}$p<0.05; $^{***}$p<0.01

# D
# Author contributions

## Chapter 1

I, Nico Gradwohl, hereby declare that I am the sole author of the introduction to this thesis (Chapter 1).

## Chapter 2

Gradwohl, N., Gaissmaier, & Neth, H. (2021). An ideal match? Learning to distribute choices under undisclosed competition with limited social information.

| | |
|---|---|
| Idea and study design: | NG, HN |
| Data collection: | NG |
| Data analysis: | NG |
| Preparation of first manuscript: | NG |
| Critical revisions: | NG, WG, HN |

## Chapter 3

Gradwohl, N., Neth, H., & Gaissmaier, W. (2021). Effects of providing and receiving recommendations on exploration behavior in spatially correlated bandits.

| | |
|---|---|
| Idea and study design: | NG, WG, HN |
| Data collection: | NG |
| Data analysis: | NG |
| Preparation of first manuscript: | NG |
| Critical revisions: | NG, WG, HN |

## Chapter 4

Gradwohl, N., Neth, H., & Gaissmaier, W. (2021). Moral reasoning in the wild: Triage decisions in the context of the COVID-19 pandemic reflect integration of information beyond a concern for the greater good.

| | |
|---|---|
| Idea and study design: | NG, WG, HN |
| Data collection: | NG |
| Data analysis: | NG |
| Preparation of first manuscript: | NG |
| Critical revisions: | NG, WG, HN |

Konstanz, March 2021

Nico Gradwohl

# Bibliography

Acerbi, A., Tennie, C., & Mesoudi, A. (2016). Social learning solves the problem of narrow-peaked search landscapes: Experimental evidence in humans. *Royal Society Open Science*, *3*(9), 160215. https://doi.org/10.1098/rsos.160215

Acerbi, A., Tennie, C., & Nunn, C. L. (2011). Modeling imitation and emulation in constrained search spaces. *Learning and Behavior*, *39*(2), 104–114. https://doi.org/10.3758/s13420-010-0009-z

Anderson, C. J. (2003). The psychology of doing nothing: Forms of decision avoidance result from reason and emotion. *Psychological Bulletin*, *129*(1), 139–167. https://doi.org/10.1037/0033-2909.129.1.139

Andreoni, J., & Bernheim, B. D. (2009). Social image and the 50–50 norm: A theoretical and experimental analysis of audience effects. *Econometrica*, *77*(5), 1607–1636. https://doi.org/10.3982/ECTA7384

Aplin, L. M. (2019). Culture and cultural evolution in birds: A review of the evidence. *Animal Behaviour*, *147*, 179–187. https://doi.org/10.1016/j.anbehav.2018.05.001

Aplin, L. M., Farine, D. R., Morand-Ferron, J., Cockburn, A., Thornton, A., & Sheldon, B. C. (2015). Experimentally induced innovations lead to persistent culture via conformity in wild birds. *Nature*, *518*(7540), 538–541. https://doi.org/10.1038/nature13998

Arechar, A. A., Kraft-Todd, G. T., & Rand, D. G. (2017). Turking overtime: How participant characteristics and behavior vary over time and day on Amazon Mechanical Turk. *Journal of the Economic Science Association*, *3*(1), 1–11. https://doi.org/10.1007/s40881-017-0035-0

Arechar, A. A., & Rand, D. (2020). Turking in the time of COVID. https://doi.org/10.31234/osf.io/vktqu

Arkes, H. R., Dawes, R. M., & Christensen, C. (1986). Factors influencing the use of a decision rule in a probabilistic task. *Organizational Behavior and Human Decision Processes*, *37*(1), 93–110. https://doi.org/10.1016/0749-5978(86)90046-4

Auer, P., Cesa-Bianchi, N., & Fischer, P. (2002). Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, *47*(2), 235–256. https://doi.org/10.1023/A:1013689704352

Avrahami, J., Güth, W., & Kareev, Y. (2005). Games of competition in a stochastic environment. *Theory and Decision*, *59*(4), 255–294. https://doi.org/10.1007/s11238-005-2883-y

Awad, E., Dsouza, S., Kim, R., Schulz, J., Henrich, J., Shariff, A., Bonnefon, J.-F., & Rahwan, I. (2018). The moral machine experiment. *Nature*, *563*(7729), 59–64. https://doi.org/10.1038/s41586-018-0637-6

Axelrod, R., & Hamilton, W. D. (1981). The evolution of cooperation. *Science*, *211*(4489), 1390–1396. https://doi.org/10.1126/science.7466396

Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, *59*(4), 390–412. https://doi.org/10.1016/j.jml.2007.12.005

Babor, T. F., Higgins-Biddle, J. C., Saunders, J. B., Monteiro, M. G., et al. (2001). *The alcohol use disorders identification test.* World Health Organization Geneva.

Balliet, D., Mulder, L. B., & Van Lange, P. A. M. (2011). Reward, punishment, and cooperation: A meta-analysis. *Psychological Bulletin*, *137*(4), 594–615. https://doi.org/10.1037/a0023489

Balliet, D., Tybur, J. M., & Van Lange, P. A. M. (2017). Functional interdependence theory: An evolutionary account of social situations. *Personality and Social Psychology Review*, *21*(4), 361–388. https://doi.org/10.1177/1088868316657965

Balliet, D., Wu, J., & De Dreu, C. K. (2014). Ingroup favoritism in cooperation: A meta-analysis. *Psychological Bulletin*, *140*(6), 1556–1581. https://doi.org/10.1037/a0037737

Bardsley, N. (2000). Control without deception: Individual behaviour in free-riding experiments revisited. *Experimental Economics*, *3*(3), 215–240. https://doi.org/10.1023/A:1011420500828

Baron, J., & Ritov, I. (1994). Reference points and omission bias. *Organizational Behavior and Human Decision Processes*, *59*(3), 475–498. https://doi.org/10.1006/obhd.1994.1070

Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, *68*(3), 255–278. https://doi.org/10.1016/j.jml.2012.11.001

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, *67*(1), 1–48. https://doi.org/10.18637/jss.v067.i01

Batson, C. D., Thompson, E. R., Seuferling, G., Whitney, H., & Strongman, J. A. (1999). Moral hypocrisy: Appearing moral to oneself without being so. *Journal of Personality and Social Psychology*, *77*(3), 525–537. https://doi.org/10.1037/0022-3514.77.3.525

Bauman, C. W., McGraw, A. P., Bartels, D. M., & Warren, C. (2014). Revisiting external validity: Concerns about trolley problems and other sacrificial dilemmas in moral psychology. *Social and Personality Psychology Compass*, *8*(9), 536–554. https://doi.org/10.1111/spc3.12131

Benjamin, D., & Budescu, D. V. (2015). Advice from experience: Communicating incomplete information incompletely. *Journal of Behavioral Decision Making*, *28*(1), 36–49. https://doi.org/10.1002/bdm.1825

Bentham, J. (1983). *The collected works of Jeremy Bentham: Deontology, together with a table of the springs of action; and the article on utilitarianism* [Original work published 1789]. Oxford University Press.

Berger-Tal, O., Nathan, J., Meron, E., & Saltz, D. (2014). The exploration-exploitation dilemma: A multidisciplinary framework. *PloS ONE*, *9*(4), e95693. https://doi.org/10.1371/journal.pone.0095693

Bernhard, H., Fischbacher, U., & Fehr, E. (2006). Parochial altruism in humans. *Nature*, *442*(7105), 912–915.

Betsch, C., Wieler, L. H., & Habersaat, K. (2020). Monitoring behavioural insights related to COVID-19. *The Lancet*, *395*(10232), 1255–1256. https://doi.org/10.1016/S0140-6736(20)30729-7

Biddison, L. D., Berkowitz, K. A., Courtney, B., De Jong, C. M. J., Devereaux, A. V., Kissoon, N., Roxland, B. E., Sprung, C. L., Dichter, J. R., Christian, M. D., et al. (2014). Ethical considerations: Care of the critically ill and injured during pandemics and disasters: CHEST consensus statement. *Chest*, *146*(4), e145S–e155S. https://doi.org/10.1378/chest.14-0742

Biele, G., Rieskamp, J., & Gonzalez, R. (2009). Computational models for the combination of advice and individual learning. *Cognitive Science*, *33*(2), 206–242. https://doi.org/10.1111/j.1551-6709.2009.01010.x

Biele, G., Rieskamp, J., Krugel, L. K., & Heekeren, H. R. (2011). The neural basis of following advice. *PLoS Biology*, *9*(6). https://doi.org/10.1371/journal.pbio.1001089

Bikhchandani, S., Hirshleifer, D., & Welch, I. (1998). Learning from the behavior of others: Conformity, fads, and informational cascades. *Journal of Economic Perspectives*, *12*(3), 151–170. https://doi.org/10.1257/jep.12.3.151

Bilancini, E., Boncinelli, L., Capraro, V., Celadin, T., & Di Paolo, R. (2020). "Do the right thing" for whom? An experiment on ingroup favouritism, group assorting and moral suasion. *Judgment and Decision Making*, *15*(2), 182–192. https://sjdm.org/journal/19/191113/jdm191113.pdf

Black, J. M. (1988). Preflight signalling in swans: A mechanism for group cohesion and flock formation. *Ethology*, *79*(2), 143–157. https://doi.org/10.1111/j.1439-0310.1988.tb00707.x

Blunden, H., & Gino, F. (2018). How the other half thinks. In E. L. MacGeorge & L. M. Van Swol (Eds.), *The Oxford handbook of advice* (pp. 43–67). https://doi.org/10.1093/oxfordhb/9780190630188.013.3

Bonaccio, S., & Dalal, R. S. (2006). Advice taking and decision-making: An integrative literature review, and implications for the organizational sciences. *Organizational Behavior and Human Decision Processes*, *101*(2), 127–151. https://doi.org/10.1016/j.obhdp.2006.07.001

Bonawitz, E., Shafto, P., Gweon, H., Goodman, N. D., Spelke, E., & Schulz, L. (2011). The double-edged sword of pedagogy: Instruction limits spontaneous exploration and discovery. *Cognition*, *120*(3), 322–330. https://doi.org/10.1016/j.cognition.2010.10.001

Boorman, E. D., O'Doherty, J. P., Adolphs, R., & Rangel, A. (2013). The behavioral and neural mechanisms underlying the tracking of expertise. *Neuron*, *80*(6), 1558–1571. https://doi.org/10.1016/j.neuron.2013.10.024

Borsari, B., & Carey, K. B. (2003). Descriptive and injunctive norms in college drinking: A meta-analytic integration. *Journal of Studies on Alcohol*, *64*(3), 331–341. https://doi.org/10.15288/jsa.2003.64.331

Bostyn, D. H., Sevenhant, S., & Roets, A. (2018). Of mice, men, and trolleys: Hypothetical judgment versus real-life behavior in trolley-style moral dilemmas. *Psychological Science*, *29*(7), 1084–1093. https://doi.org/10.1177/0956797617752640

Boyd, R., Gintis, H., Bowles, S., & Richerson, P. J. (2003). The evolution of altruistic punishment. *Proceedings of the National Academy of Sciences*, *100*(6), 3531–3535. https://doi.org/10.1073/pnas.0630443100

Boyd, R., & Richerson, P. J. (1985). *Culture and the evolutionary process*. University of Chicago press.

Boyd, R., & Richerson, P. J. (1995). Why does culture increase human adaptability? *Ethology and Sociobiology*, *16*(2), 125–143. https://doi.org/10.1016/0162-3095(94)00073-G

Briegel, H. J., & De las Cuevas, G. (2012). Projective simulation for artificial intelligence. *Scientific Reports*, *2*(400), 1–16. https://doi.org/10.1038/srep00400

Bruine de Bruin, W., Parker, A. M., Galesic, M., & Vardavas, R. (2019). Reports of social circles' and own vaccination behavior: A national longitudinal survey. *Health Psychology*, *38*(11), 975–983. https://doi.org/10.1037/hea0000771

Buckwalter, W., & Peterson, A. (2020). Public attitudes toward allocating scarce resources in the COVID-19 pandemic. *PLoS ONE*, *15*(11), 1–20. https://doi.org/10.1371/journal.pone.0240651

Caldwell, C. A., Atkinson, M., Blakey, K. H., Dunstone, J., Kean, D., Mackintosh, G., Renner, E., & Wilks, C. E. H. (2020). Experimental assessment of capacities for cumulative culture: Review and evaluation of methods. *Wiley Interdisciplinary Reviews: Cognitive Science*, *11*(1), 1–15. https://doi.org/10.1002/wcs.1516

Caldwell, C. A., & Millen, A. E. (2009). Social learning mechanisms and cumulative cultural evolution: Is imitation necessary? *Psychological Science*, *20*(12), 1478–1483. https://doi.org/10.1111/j.1467-9280.2009.02469.x

Camerer, C. (2003). *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press.

Camerer, C., & Mobbs, D. (2017). Differences in behavior and brain activity during hypothetical and real choices. *Trends in Cognitive Sciences*, *21*(1), 46–56. https://doi.org/10.1016/j.tics.2016.11.001

Carnes, N. C., Lickel, B., & Janoff-Bulman, R. (2015). Shared perceptions: Morality is embedded in social contexts. *Personality and Social Psychology Bulletin*, *41*(3), 351–362. https://doi.org/10.1177/0146167214566187

Çelen, B., Kariv, S., & Schotter, A. (2010). An experimental test of advice and social learning. *Management Science*, *56*(10), 1687–1701. https://doi.org/10.1287/mnsc.1100.1228

Chari, V. V., & Kehoe, P. J. (2004). Financial crises as herds: Overturning the critiques. *Journal of Economic Theory*, *119*(1), 128–150. https://doi.org/10.1016/S0022-0531(03)00225-4

Chen, D. L., Schonger, M., & Wickens, C. (2016). oTree—An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, *9*, 88–97. https://doi.org/10.1016/j.jbef.2015.12.001

Christensen, J. F., & Gomila, A. (2012). Moral dilemmas in cognitive neuroscience of moral decision-making: A principled review. *Neuroscience & Biobehavioral Reviews*, *36*(4), 1249–1264. https://doi.org/10.1016/j.neubiorev.2012.02.008

Cialdini, R. B., & Trost, M. R. (1998). Social influence: Social norms, conformity and compliance. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *The handbook of social psychology* (pp. 151–192). McGraw-Hill.

Cialdini, R. B., & Goldstein, N. J. (2004). Social influence: Compliance and conformity. *Annual Review of Psychology*, *55*(1), 591–621. https://doi.org/10.1146/annurev.psych.55.090902.142015

Cikara, M., Farnsworth, R. A., Harris, L. T., & Fiske, S. T. (2010). On the wrong side of the trolley track: Neural correlates of relative social valuation. *Social Cognitive and Affective Neuroscience*, *5*(4), 404–413. https://doi.org/10.1093/scan/nsq011

Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *362*(1481), 933–942. https://doi.org/10.1098/rstb.2007.2098

Conradt, L., & Roper, T. J. (2005). Consensus decision making in animals. *Trends in Ecology and Evolution*, *20*(8), 449–456. https://doi.org/10.1016/j.tree.2005.05.008

Conway, P., Goldstein-Greenwood, J., Polacek, D., & Greene, J. D. (2018). Sacrificial utilitarian judgments do reflect concern for the greater good: Clarification via process dissociation and the judgments of philosophers. *Cognition*, *179*, 241–265. https://doi.org/10.1016/j.cognition.2018.04.018

Couzin, I. D., & Franks, N. R. (2003). Self-organized lane formation and optimized traffic flow in army ants. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, *270*(1511), 139–146. https://doi.org/10.1098/rspb.2002.2210

Couzin, I. D., Ioannou, C. C., Demirel, G., Gross, T., Torney, C. J., Hartnett, A., Conradt, L., Levin, S. A., & Leonard, N. E. (2011). Uninformed individuals promote democratic consensus in animal groups. *Science*, *334*(6062), 1578–1580. https://doi.org/10.1126/science.1210280

Couzin, I. D. (2009). Collective cognition in animal groups. *Trends in Cognitive Sciences*, *13*(1), 36–43. https://doi.org/10.1016/j.tics.2008.10.002

Critchfield, T. S., & Atteberry, T. (2003). Temporal discounting predicts individual competitive success in a human analogue of group foraging. *Behavioural Processes*, *64*(3), 315–331. https://doi.org/10.1016/S0376-6357(03)00129-3

Curşeu, P. L., Jansen, R. J. G., & Chappin, M. M. H. (2013). Decision rules and group rationality: Cognitive gain or standstill? *PLoS ONE*, *8*(2), e56454. https://doi.org/10.1371/journal.pone.0056454

Cushman, F., Young, L., & Hauser, M. (2006). The role of conscious reasoning and intuition in moral judgment: Testing three principles of harm. *Psychological Science*, *17*(12), 1082–1089. https://doi.org/10.1111/j.1467-9280.2006.01834.x

Dana, J., & Cain, D. M. (2015). Advice versus choice. *Current Opinion in Psychology*, *6*, 173–176. https://doi.org/10.1016/j.copsyc.2015.08.019

Dana, J., Weber, R. A., & Kuang, J. X. (2007). Exploiting moral wiggle room: Experiments demonstrating an illusory preference for fairness. *Economic Theory*, *33*(1), 67–80. https://doi.org/10.1007/s00199-006-0153-z

Dawes, R. M. (1980). Social dilemmas. *Annual Review of Psychology*, *31*(1), 169–193. https://doi.org/10.1146/annurev.ps.31.020180.001125

Dean, L. G., Vale, G. L., Laland, K. N., Flynn, E., & Kendal, R. L. (2014). Human cumulative culture: A comparative perspective. *Biological Reviews*, *89*(2), 284–301. https://doi.org/10.1111/brv.12053

Decety, J., & Cacioppo, S. (2012). The speed of morality: A high-density electrical neuroimaging study. *Journal of Neurophysiology*, *108*(11), 3068–3072. https://doi.org/10.1152/jn.00473.2012

DeLosh, E. L., Busemeyer, J. R., & McDaniel, M. A. (1997). Extrapolation: The sine qua non for abstraction in function learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *23*(4), 968. https://doi.org/10.1037/0278-7393.23.4.968

de Quervain, D. J.-F., Fischbacher, U., Treyer, V., Schellhammer, M., Schnyder, U., Buck, A., & Fehr, E. (2004). The neural basis of altruistic punishment. *Science*, *305*(5688), 1254–1258. https://doi.org/10.1126/science.1100735

Derex, M., & Boyd, R. (2016). Partial connectivity increases cultural accumulation within groups. *Proceedings of the National Academy of Sciences of the United States of America*, *113*(11), 2982–2987. https://doi.org/10.1073/pnas.1518798113

Derex, M., Godelle, B., & Raymond, M. (2013). Social learners require process information to outperform individual learners. *Evolution*, *67*(3), 688–697. https://doi.org/10.1111/j.1558-5646.2012.01804.x

DeScioli, P., Bruening, R., & Kurzban, R. (2011). The omission effect in moral cognition: Toward a functional explanation. *Evolution and Human Behavior*, *32*(3), 204–215. https://doi.org/10.1016/j.evolhumbehav.2011.01.003

Deutsch, M., & Gerard, H. B. (1955). A study of normative and informational social influences upon individual judgment. *The Journal of Abnormal and Social Psychology*, *51*(3), 629–636. https://doi.org/10.1037/h0046408

Diaconescu, A. O., Mathys, C., Weber, L. A. E., Daunizeau, J., Kasper, L., Lomakina, E. I., Fehr, E., & Stephan, K. E. (2014). Inferring on the intentions of others by hierarchical Bayesian learning. *PLoS Computational Biology*, *10*(9). https://doi.org/10.1371/journal.pcbi.1003810

Dijker, A. J. M., Nelissen, R. M. A., & Stijnen, M. M. N. (2013). Framing posthumous organ donation in terms of reciprocity: What are the emotional consequences? *Basic and Applied Social Psychology*, *35*(3), 256–264. https://doi.org/10.1080/01973533.2013.785401

Disma, G., Sokolowski, M. B. C., & Tonneau, F. (2011). Children's competition in a natural setting: Evidence for the ideal free distribution. *Evolution and Human Behavior*, *32*(6), 373–379. https://doi.org/10.1016/j.evolhumbehav.2010.11.007

Dunwoody, P. T. (2009). Theories of truth as assessment criteria in judgment and decision making. *Judgment and Decision Making*, *4*(2), 116. http://journal.sjdm.org/ccd/ccd.pdf

Dwenger, N., Kübler, D., & Weizsäcker, G. (2012). *Flipping a coin: Theory and evidence* (Working Paper No. 2012–14). Max Planck Institute for Tax Law and Public Finance. Munich, Germany.

Edwards, W. (1961). Probability learning in 1000 trials. *Journal of Experimental Psychology*, *62*(4), 385–394. https://doi.org/10.1037/h0041970

Efferson, C., Lalive, R., Richerson, P. J., McElreath, R., & Lubell, M. (2008). Conformists and mavericks: The empirics of frequency-dependent cultural transmission. *Evolution and Human Behavior*, *29*(1), 56–64. https://doi.org/10.1016/j.evolhumbehav.2007.08.003

Emanuel, E. J., Persad, G., Upshur, R., Thome, B., Parker, M., Glickman, A., Zhang, C., Boyle, C., Smith, M., & Phillips, J. P. (2020). Fair allocation of scarce medical resources in the time of Covid-19. *New England Journal of Medicine*, *382*(21), 2049–2055. https://doi.org/10.1056/NEJMsb2005114

Engel, C. (2011). Dictator games: A meta study. *Experimental Economics*, *14*(4), 583–610. https://doi.org/10.1007/s10683-011-9283-7

Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review*, *112*(4), 912–931. https://doi.org/10.1037/0033-295X.112.4.912

Erev, I., & Rapoport, A. (1998). Coordination, "magic," and reinforcement learning in a market entry game. *Games and Economic Behavior*, *23*(2), 146–175. https://doi.org/10.1006/game.1997.0619

Estes, W. K. (1976). The cognitive side of probability learning. *Psychological Review*, *83*(1), 37–64. https://doi.org/10.1037/0033-295X.83.1.37

Everett, J. A. C., Faber, N. S., & Crockett, M. (2015). Preferences and beliefs in ingroup favoritism. *Frontiers in Behavioral Neuroscience*, *9*, 15. https://doi.org/10.3389/fnbeh.2015.00015

Everett, J. A. C., & Kahane, G. (2020). Switching tracks? Towards a multidimensional model of utilitarian psychology. *Trends in Cognitive Sciences*, *24*(2), 124–134. https://doi.org/10.1016/j.tics.2019.11.012

Falk, A., Becker, A., Dohmen, T. J., Huffman, D., & Sunde, U. (2016). The preference survey module: A validated instrument for measuring risk, time, and social preferences.

Falk, A., & Hermle, J. (2018). Relationship of gender differences in preferences to economic development and gender equality. *Science*, *362*(6412). https://doi.org/10.1126/science.aas9899

Fallucchi, F., Faravelli, M., & Quercia, S. (2020). Fair allocation of scarce medical resources in the time of COVID-19: What do people think? *Journal of Medical Ethics*. https://doi.org/10.1136/medethics-2020-106524

Fantino, E., & Esfandiari, A. (2002). Probability matching: Encouraging optimal responding in humans. *Canadian Journal of Experimental Psychology*, *56*(1), 58–63. https://doi.org/10.1037/h0087385

Feher da Silva, C., Victorino, C. G., Caticha, N., & Baldo, M. V. C. (2017). Exploration and recency as the main proximate causes of probability matching: A reinforcement learning analysis. *Scientific Reports*, *7*, 1–23. https://doi.org/10.1038/s41598-017-15587-z

Fehr, E., & Fischbacher, U. (2004). Third-party punishment and social norms. *Evolution and Human Behavior*, *25*(2), 63–87. https://doi.org/10.1016/S1090-5138(04)00005-4

Fehr, E., & Gächter, S. (2002). Altruistic punishment in humans. *Nature*, *415*(6868), 137–140. https://doi.org/10.1038/415137a

FeldmanHall, O., Mobbs, D., Evans, D., Hiscox, L., Navrady, L., & Dalgleish, T. (2012). What we say and what we do: The relationship between real and hypothetical moral choices. *Cognition*, *123*(3), 434–441. https://doi.org/10.1016/j.cognition.2012.02.001

FeldmanHall, O., Sokol-Hessner, P., Van Bavel, J. J., & Phelps, E. A. (2014). Fairness violations elicit greater punishment on behalf of another than for oneself. *Nature Communications*, *5*(1), 1–6. https://doi.org/10.1038/ncomms6306

Fiedler, K., Hütter, M., Schott, M., & Kutzner, F. (2019). Metacognitive myopia and the overutilization of misleading advice. *Journal of Behavioral Decision Making*, *32*(3), 317–333. https://doi.org/10.1002/bdm.2109

Foot, P. (1967). The problem of abortion and the doctrine of double effect. *Oxford Review*, *5*, 5–15.

Fowler, J. H. (2005). Altruistic punishment and the origin of cooperation. *Proceedings of the National Academy of Sciences*, *102*(19), 7047–7049. https://doi.org/10.1073/pnas.0500938102

Fox, J., & Weisberg, S. (2019). *An R companion to applied regression* (Third). Sage. https://socialsciences.mcmaster.ca/jfox/Books/Companion/

Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives*, *19*(4), 25–42. https://doi.org/10.1257/089533005775196732

Fretwell, S. D., & Lucas, H. L. (1969). On territorial behavior and other factors influencing habitat distribution in birds. *Acta Biotheoretica*, *19*(1), 16–36. https://doi.org/10.1007/BF01601953

Frith, C. D., & Frith, U. (2012). Mechanisms of Social Cognition. *Annual Review of Psychology*, *63*(1), 287–313. https://doi.org/10.1146/annurev-psych-120710-100449

Gaissmaier, W., & Schooler, L. J. (2008). The smart potential behind probability matching. *Cognition*, *109*(3), 416–422. https://doi.org/10.1016/j.cognition.2008.09.007

Gallistel, C. R. (1990). *Learning, development, and conceptual change. The organization of learning.* The MIT Press.

Gardner, P. H., & Berry, D. C. (1995). The effect of different forms of advice on the control of a simulated complex system. *Applied Cognitive Psychology*, *9*(7), 55–79. https://doi.org/10.1002/acp.2350090706

Gelman, A., & Hill, J. (2007). *Data analysis using regression and multilevel/hierarchical models.* Cambridge University Press.

Gershman, S. J. (2019). Uncertainty and exploration. *Decision*, *6*(3), 277–286. https://doi.org/10.1037/dec0000101

Gershman, S. J., & Daw, N. D. (2017). Reinforcement learning and episodic memory in humans and animals: An integrative framework. *Annual Review of Psychology*, *68*(1), 101–128. https://doi.org/10.1146/annurev-psych-122414-033625

Giese, H., Stok, F. M., & Renner, B. (2019). Perceiving college peers' alcohol consumption: Temporal patterns and individual differences in overestimation. *Psychology & Health*, *34*(2), 147–161. https://doi.org/10.1080/08870446.2018.1514118

Gigerenzer, G., Todd, P. M., & the ABC Research Group. (1999). *Simple heuristics that make us smart.* Oxford University Press.

Gigerenzer, G. (1996). Rationality: Why social context matters. In P. B. Baltes & U. M. Staudinger (Eds.), *Interactive minds: Life-span perspectives on the social foundation of cognition* (pp. 319–346). Cambridge University Press.

Gigerenzer, G. (2010). Moral satisficing: Rethinking moral behavior as bounded rationality. *Topics in Cognitive Science*, *2*(3), 528–554. https://doi.org/10.1111/j.1756-8765.2010.01094.x

Gigerenzer, G. (2020). How to explain behavior? *Topics in Cognitive Science*, *12*(4), 1363–1381. https://doi.org/10.1111/tops.12480

Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology*, *62*, 451–482. https://doi.org/10.1146/annurev-psych-120709-145346

Giraldeau, L. A., Valone, T. J., & Templeton, J. J. (2002). Potential disadvantages of using socially acquired information. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *357*(1427), 1559–1566. https://doi.org/10.1098/rstb.2002.1065

Goldstein, N. J., Cialdini, R. B., & Griskevicius, V. (2008). A room with a viewpoint: Using social norms to motivate environmental conservation in hotels. *Journal of Consumer Research*, *35*(3), 472–482. https://doi.org/10.1086/586910

Goldstone, R. L., & Ashpole, B. C. (2004). Human foraging behavior in a virtual environment. *Psychonomic Bulletin & Review*, *11*(3), 508–514. https://doi.org/10.3758/BF03196603

Goldstone, R. L., Ashpole, B. C., & Roberts, M. E. (2005). Knowledge of resources and competitors in human foraging. *Psychonomic Bulletin & Review*, *12*(1), 81–87. https://doi.org/10.3758/BF03196350

Goldstone, R. L., & Gureckis, T. M. (2009). Collective behavior. *Topics in Cognitive Science*, *1*(3), 412–438. https://doi.org/10.1111/j.1756-8765.2009.01038.x

Greene, J. D., & Haidt, J. (2002). How (and where) does moral judgment work? *Trends in Cognitive Sciences*, *6*(12), 517–523. https://doi.org/10.1016/S1364-6613(02)02011-9

Guerin, B. (1986). Mere presence effects in humans: A review. *Journal of Experimental Social Psychology*, *22*(1), 38–77. https://doi.org/10.1016/0022-1031(86)90040-5

Gurmankin, A. D., Baron, J., Hershey, J. C., & Ubel, P. A. (2002). The role of physicians' recommendations in medical treatment decisions. *Medical Decision Making*, *22*(3), 262–271. https://doi.org/10.1177/0272989X0202200314

Haidt, J. (2001). The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review*, *108*(4), 814–834. https://doi.org/10.1037/0033-295X.108.4.814

Harries, C., Yaniv, I., & Harvey, N. (2004). Combining advice: The weight of a dissenting opinion in the consensus. *Journal of Behavioral Decision Making*, *17*(5), 333–348. https://doi.org/10.1002/bdm.474

Harvey, N., & Fischer, I. (1997). Taking advice: Accepting help, improving judgment, and sharing responsibility. *Organizational Behavior and Human Decision Processes*, *70*(2), 117–133. https://doi.org/10.1006/obhd.1997.2697

Helbing, D., Buzna, L., Johansson, A., & Werner, T. (2005). Self-organized pedestrian crowd dynamics: Experiments, simulations, and design solutions. *Transportation Science*, *39*(1), 1–24. https://doi.org/10.1287/trsc.1040.0108

Helbing, D., Molnár, P., Farkas, I. J., & Bolay, K. (2001). Self-organizing pedestrian movement. *Environment and Planning B: Planning and Design*, *28*(3), 361–383. https://doi.org/10.1068/b2697

Henrich, J., Ensminger, J., McElreath, R., Barr, A., Barrett, C., Bolyanatz, A., Cardenas, J. C., Gurven, M., Gwako, E., Henrich, N., Lesorogol, C., Marlowe, F., Tracer, D., & Ziker, J. (2010). Markets, religion, community size, and the evolution of fairness and punishment. *Science*, *327*(5972), 1480–1484. https://doi.org/10.1126/science.1182238

Henrich, J., & Gil-White, F. J. (2001). The evolution of prestige: Freely conferred deference as a mechanism for enhancing the benefits of cultural transmission. *Evolution and Human Behavior*, *22*(3), 165–196. https://doi.org/10.1016/S1090-5138(00)00071-4

Herrnstein, R. J., & Prelec, D. (1991). Melioration: A theory of distributed choice. *Journal of Economic Perspectives*, *5*(3), 137–156. https://doi.org/10.1257/jep.5.3.137

Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, *15*(8), 534–539. https://doi.org/10.1111/j.0956-7976.2004.00715.x

Hertwig, R., & Erev, I. (2009). The description–experience gap in risky choice. *Trends in Cognitive Sciences*, *13*(12), 517–523. https://doi.org/10.1016/j.tics.2009.09.004

Hertwig, R., & Hoffrage, U. (2012). Simple Heuristics: The Foundations of Adaptive Social Behavior. In R. Hertwig, U. Hoffrage, & The ABC Research Group (Eds.), *Simple heuristics in a social world* (pp. 3–36). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780195388435.003.0001

Hertwig, R., Hoffrage, U., & Research Group, A. (Eds.). (2012). *Simple Heuristics in a Social World*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780195388435.001.0001

Heylighen, F. (2016). Stigmergy as a universal coordination mechanism II: Varieties and evolution. *Cognitive Systems Research*, *38*, 50–59. https://doi.org/10.1016/j.cogsys.2015.12.007

Hick, J. L., Hanfling, D., Wynia, M. K., & Pavia, A. T. (2020). Duty to plan: Health care, crisis standards of care, and novel coronavirus SARS-CoV-2. *NAM Perspectives*. https://doi.org/10.31478/202003

Hills, T. T., Todd, P. M., Lazer, D., Redish, A. D., & Couzin, I. D. (2015). Exploration versus exploitation in space, mind, and society. *Trends in Cognitive Sciences*, *19*(1), 46–54. https://doi.org/10.1016/j.tics.2014.10.004

Hoffman, P. J. (1960). The paramorphic representation of clinical judgment. *Psychological Bulletin*, *57*(2), 116–131. https://doi.org/10.1037/h0047807

Holyoak, K. J., & Powell, D. (2016). Deontological coherence: A framework for commonsense moral reasoning. *Psychological Bulletin*, *142*(11), 1179–1203. https://doi.org/10.1037/bul0000075

Hütter, M., & Fiedler, K. (2019). Advice taking under uncertainty: The impact of genuine advice versus arbitrary anchors on judgment. *Journal of Experimental Social Psychology*, *85*, 103829. https://doi.org/10.1016/j.jesp.2019.103829

Huyck, J. B. V., Battalio, R. C., & Beil, R. O. (1990). Tacit coordination games, strategic uncertainty, and coordination failure. *The American Economic Review*, *80*(1), 234–248. http://www.jstor.org/stable/2006745

IRS. (2019). *Comprehensive Taxpayer Attitude Survey (CTAS) 2019* (Executive Report). Retrieved February 16, 2021, from https://www.irs.gov/pub/irs-pdf/p5296.pdf

Jaeger, T. F. (2008). Categorical data analysis: Away from ANOVAs (transformation or not) and towards logit mixed models. *Journal of Memory and Language*, *59*(4), 434–446. https://doi.org/10.1016/j.jml.2007.11.007

Jaques, A. E. (2019). Why the moral machine is a monster. *University of Miami School of Law*, *10*.

Jaziri, R., & Alnahdi, S. (2020). Choosing which COVID-19 patient to save? the ethical triage and rationing dilemma. *Ethics, Medicine and Public Health*, *15*, 100570. https://doi.org/10.1016/j.jemep.2020.100570

Jiménez, Á. V., & Mesoudi, A. (2019). Prestige-biased social learning: Current evidence and outstanding questions. *Palgrave Communications*, *5*(1), 1–12. https://doi.org/10.1057/s41599-019-0228-7

Jin, L., Huang, Y., Liang, Y., & Zhang, Q. (2021). Who gets the ventilator? Moral decision making regarding medical resource allocation in a pandemic. *Journal of the Association for Consumer Research*, *6*(1), 159–167. https://doi.org/10.1086/711734

Jonas, E., & Frey, D. (2003). Information search and presentation in advisor–client interactions. *Organizational Behavior and Human Decision Processes*, *91*(2), 154–168. https://doi.org/10.1016/S0749-5978(03)00059-1

Kahane, G., Everett, J. A., Earp, B. D., Caviola, L., Faber, N. S., Crockett, M. J., & Savulescu, J. (2018). Beyond sacrificial harm: A two-dimensional model of utilitarian psychology. *Psychological Review*, *125*(2), 131. https://doi.org/10.1037/rev0000093

Kameda, T., & Hastie, R. (2015). Herd behavior. *Emerging Trends in the Social and Behavioral Sciences: An Interdisciplinary, Searchable, and Linkable Resource*, 1–14. https://doi.org/10.1002/9781118900772.etrds0157

Kameda, T., & Nakanishi, D. (2002). Cost–benefit analysis of social/cultural learning in a nonstationary uncertain environment. *Evolution and Human Behavior*, *23*(5), 373–393. https://doi.org/10.1016/S1090-5138(02)00101-0

Kareev, Y., Avrahami, J., Goldzweig, A., Hadar, D., & Klein, S. (2019). Myopic reactivity in repeated-choice: The case of competing with others for a scarce, unevenly distributed resource. *Journal of Experimental Psychology: General*, *149*, 774–789. https://doi.org/10.1037/xge0000676

KC, R. P., Kunter, M., & Mak, V. (2018). The influence of a competition on noncompetitors. *Proceedings of the National Academy of Sciences*, *115*(11), 2716–2721. https://doi.org/10.1073/pnas.1717301115

Kelley, H. H., & Thibaut, J. W. (1978). *Interpersonal relations: A theory of interdependence*. New York: Wiley.

Kelley, H. H., Thibaut, J. W., Radloff, R., & Mundy, D. (1962). The development of cooperation in the “minimal social situation”. *Psychological Monographs: General and Applied*, *76*(1–19), 1. https://doi.org/10.1037/h0093819

Kendal, R. L., Boogert, N. J., Rendell, L., Laland, K. N., Webster, M., & Jones, P. L. (2018). Social learning strategies: Bridge-building between fields. *Trends in Cognitive Sciences*, *22*(7), 651–665. https://doi.org/10.1016/j.tics.2018.04.003

Kennedy, M., & Gray, R. D. (1993). Can ecological theory predict the distribution of foraging animals? A critical analysis of experiments on the Ideal Free Distribution. *Oikos*, *68*(1), 158. https://doi.org/10.2307/3545322

Keren, G., & Teigen, K. H. (2010). Decisions by coin toss: Inappropriate but fair. *Judgment and Decision Making*, *5*(2), 83–101. http://journal.sjdm.org/10/10203/jdm10203.pdf

Kneer, M., & Hannikainen, I. R. (2021). Trolleys, triage and Covid-19: The role of psychological realism in sacrificial dilemmas. *Cognition and Emotion*, *0*(0), 1–17. https://doi.org/10.1080/02699931.2021.1964940

Knight, F. H. (1921). *Risk, uncertainty and profit*. University of Chicago Press.

Koehler, D. J., & James, G. (2009). Probability matching in choice under uncertainty: Intuition versus deliberation. *Cognition*, *113*(1), 123–127. https://doi.org/10.1016/j.cognition.2009.07.003

Koehler, D. J., & James, G. (2010). Probability matching and strategy availability. *Memory & Cognition*, *38*(6), 667–676. https://doi.org/10.3758/MC.38.6.667

Korn, L., Böhm, R., Meier, N. W., & Betsch, C. (2020). Vaccination as a social contract. *Proceedings of the National Academy of Sciences*, *117*(26), 14890–14899. https://doi.org/10.1073/pnas.1919666117

Kraft, J. R., & Baum, W. M. (2001). Group choice: The ideal free distribution of human social behavior. *Journal of the Experimental Analysis of Behavior*, *76*(1), 21–42. https://doi.org/10.1901/jeab.2001.76-21

Kraft, J. R., Baum, W. M., & Burge, M. J. (2002). Group choice and individual choices: Modeling human social behavior with the Ideal Free Distribution. *Behavioural Processes*, *57*(2-3), 227–240. https://doi.org/10.1016/S0376-6357(02)00016-5

Kruschke, J. K., & Johansen, M. K. (1999). A model of probabilistic category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*(5), 1083–1119. https://doi.org/10.1037/0278-7393.25.5.1083

Kubota, J. T., Li, J., Bar-David, E., Banaji, M. R., & Phelps, E. A. (2013). The price of racial bias: Intergroup negotiations in the ultimatum game. *Psychological Science*, *24*(12), 2498–2504. https://doi.org/10.1177/0956797613496435

Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest Package: Tests in Linear Mixed Effects Models. *Journal of Statistical Software*, *82*(13), 1–26. https://doi.org/10.18637/jss.v082.i13

Lai, T. L., & Robbins, H. (1985). Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, *6*(1), 4–22. https://doi.org/10.1016/0196-8858(85)90002-8

Laland, K. N. (2004). Social learning strategies. *Animal Learning & Behavior*, *32*(1), 4–14. https://doi.org/10.3758/BF03196002

Laland, K. N., Toyokawa, W., & Oudman, T. (2020). Animal learning as a source of developmental bias. *Evolution and Development*, *22*(1-2), 126–142. https://doi.org/10.1111/ede.12311

Laland, K. N., & Williams, K. (1998). Social transmission of maladaptive information in the guppy. *Behavioral Ecology*, *9*(5), 493–499. https://doi.org/10.1093/beheco/9.5.493

Lee, M. D., & Cummins, T. D. R. (2004). Evidence accumulation in decision making: Unifying the "take the best" and the "rational" models. *Psychonomic Bulletin & Review*, *11*(2), 343–352. https://doi.org/10.3758/BF03196581

Legare, C. H., & Nielsen, M. (2015). Imitation and innovation: The dual engines of cultural learning. *Trends in Cognitive Sciences*, *19*(11), 688–699. https://doi.org/10.1016/j.tics.2015.08.005

Lenth, R. (2019). emmeans: Estimated Marginal Means, aka Least-Squares Means. https://cran.r-project.org/package=emmeans

Leonhardt, J. M., Keller, L. R., & Pechmann, C. (2011). Avoiding the risk of responsibility by seeking uncertainty: Responsibility aversion and preference for indirect agency when choosing for others. *Journal of Consumer Psychology*, *21*(4), 405–413. https://doi.org/10.1016/j.jcps.2011.01.001

Littlepage, G., Robison, W., & Reddington, K. (1997). Effects of task experience and group experience on group performance, member ability, and recognition of expertise. *Organizational Behavior and Human Decision Processes*, *69*(2), 133–147. https://doi.org/10.1006/obhd.1997.2677

Lorenz, J., Rauhut, H., Schweitzer, F., & Helbing, D. (2011). How social influence can undermine the wisdom of crowd effect. *Proceedings of the National Academy of Sciences*, *108*(22), 9020–9025. https://doi.org/10.1073/pnas.1008636108

Madden, G. J., Peden, B. F., & Yamaguchi, T. (2002). Human group choice: Discrete-trial and free-operant tests of the ideal free distribution. *Journal of the Experimental Analysis of Behavior*, *78*(1), 1–15. https://doi.org/10.1901/jeab.2002.78-1

Marrow, A. J. (1969). *The practical theorist: The life and work of Kurt Lewin*. Basic Books.

Mata, A., Fiedler, K., Ferreira, M. B., & Almeida, T. (2013). Reasoning about others' reasoning. *Journal of Experimental Social Psychology*, *49*(3), 486–491. https://doi.org/10.1016/j.jesp.2013.01.010

McElreath, R., Bell, A. V., Efferson, C., Lubell, M., Richerson, P. J., & Waring, T. (2008). Beyond existence and aiming outside the laboratory: Estimating frequency-dependent and pay-off-biased social learning strategies. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *363*(1509), 3515–3528. https://doi.org/10.1098/rstb.2008.0131

McElreath, R., Lubell, M., Richerson, P. J., Waring, T. M., Baum, W., Edsten, E., Efferson, C., & Paciotti, B. (2005). Applying evolutionary models to the laboratory study of social learning. *Evolution and Human Behavior*, *26*(6), 483–508. https://doi.org/10.1016/j.evolhumbehav.2005.04.003

Mehlhorn, K., Newell, B. R., Todd, P. M., Lee, M. D., Morgan, K., Braithwaite, V. A., Hausmann, D., Fiedler, K., & Gonzalez, C. (2015). Unpacking the exploration-exploitation tradeoff: A synthesis of human and animal literatures. *Decision*, *2*(3), 191–215. https://doi.org/10.1037/dec0000033

Mesoudi, A. (2011). An experimental comparison of human social learning strategies: Payoff-biased social learning is adaptive but underused. *Evolution and Human Behavior*, *32*(5), 334–342. https://doi.org/10.1016/j.evolhumbehav.2010.12.001

Mesoudi, A. (2016). Cultural evolution: Integrating psychology, evolution and culture. *Current Opinion in Psychology*, *7*, 17–22. https://doi.org/10.1016/j.copsyc.2015.07.001

Mesoudi, A., & O'Brien, M. J. (2008a). The cultural transmission of Great Basin projectile-point technology I: An experimental simulation. *American Antiquity*, *73*(1), 3–28. https://doi.org/10.1017/S0002731600041263

Mesoudi, A., & O'Brien, M. J. (2008b). The cultural transmission of Great Basin projectile-point technology II: An agent-based computer simulation. *American Antiquity*, *73*(4), 627–644. https://doi.org/10.1017/S0002731600047338

Mesoudi, A., & Whiten, A. (2008). The multiple roles of cultural transmission experiments in understanding human cultural evolution. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *363*(1509), 3489–3501. https://doi.org/10.1098/rstb.2008.0129

Meyer, D. J., Van Huyck, J. B., Battalio, R. C., & Saving, T. R. (1992). History's role in coordinating decentralized allocation decisions. *Journal of Political Economy*, *100*(2), 292–316. https://doi.org/10.1086/261819

Mikhail, J. (2007). Universal moral grammar: Theory, evidence and the future. *Trends in Cognitive Sciences*, *11*(4), 143–152. https://doi.org/10.1016/j.tics.2006.12.007

Milinski, M., Semmann, D., & Krambeck, H.-J. (2002). Reputation helps solve the 'tragedy of the commons'. *Nature*, *415*(6870), 424–426. https://doi.org/10.1038/415424a

Mill, J. S. (1863). *Utilitarianism*. Parker, Son, Bourne.

Miton, H., & Charbonneau, M. (2018). Cumulative culture in the laboratory: Methodological and theoretical challenges. *Proceedings of the Royal Society B: Biological Sciences*, *285*(1879). https://doi.org/10.1098/rspb.2018.0677

Molleman, L., Kurvers, R. H. J. M., & van den Bos, W. (2019). Unleashing the BEAST: a brief measure of human social information use. *Evolution and Human Behavior*, *40*(5), 492–499. https://doi.org/10.1016/j.evolhumbehav.2019.06.005

Molleman, L., Van Den Berg, P., & Weissing, F. J. (2014). Consistent individual differences in human social learning strategies. *Nature Communications*, *5*, 1–9. https://doi.org/10.1038/ncomms4570

Morgan, T. J. H., Rendell, L. E., Ehn, M., Hoppitt, W., & Laland, K. N. (2012). The evolutionary basis of human social learning. *Proceedings of the Royal Society B: Biological Sciences*, *279*(1729), 653–662. https://doi.org/10.1098/rspb.2011.1172

Moussaid, M., Garnier, S., Theraulaz, G., & Helbing, D. (2009). Collective information processing and pattern formation in swarms, flocks, and crowds. *Topics in Cognitive Science*, *1*(3), 469–497. https://doi.org/10.1111/j.1756-8765.2009.01028.x

Moussaïd, M., Brighton, H., & Gaissmaier, W. (2015). The amplification of risk in experimental diffusion chains. *Proceedings of the National Academy of Sciences of the United States of America*, *112*(18), 5631–5636. https://doi.org/10.1073/pnas.1421883112

Moussaïd, M., Kämmer, J. E., Analytis, P. P., & Neth, H. (2013). Social influence and the collective dynamics of opinion formation. *PLoS ONE*, *8*(11). https://doi.org/10.1371/journal.pone.0078433

Murphy, R. O., Ackermann, K. A., & Handgraaf, M. (2011). Measuring social value orientation. *Judgment and Decision Making*, *6*(8), 771–781. journal.sjdm.org/11/m25/m25.pdf

Nash, J. (1951). Non-cooperative games. *Annals of Mathematics*, *54*(2), 286–295. https://doi.org/10.2307/1969529

Navajas, J., Heduan, F. Á., Garbulsky, G., Tagliazucchi, E., Ariely, D., & Sigman, M. (2021). Moral responses to the COVID-19 crisis. *Royal Society Open Science*, *8*(9), 210096. https://doi.org/10.1098/rsos.210096

Neth, H., Sims, C. R., & Gray, W. D. (2016). Rational task analysis: A methodology to benchmark bounded rationality. *Minds and Machines*, *26*(1-2), 125–148. https://doi.org/10.1007/s11023-015-9368-8

Newell, B. R., Koehler, D. J., James, G., Rakow, T., & van Ravenzwaaij, D. (2013). Probability matching in risky choice: The interplay of feedback and strategy availability. *Memory & Cognition*, *41*(3), 329–338. https://doi.org/10.3758/s13421-012-0268-3

Newell, B. R., & Lee, M. D. (2011). The right tool for the job? Comparing an evidence accumulation and a naive strategy selection model of decision making. *Journal of Behavioral Decision Making*, *24*(5), 456–481. https://doi.org/10.1002/bdm.703

Newell, B. R., McDonald, R. I., Brewer, M., & Hayes, B. K. (2014). The psychology of environmental decisions. *Annual Review of Environment and Resources*, *39*(1), 443–467. https://doi.org/10.1146/annurev-environ-010713-094623

Nowak, M. A., & Sigmund, K. (1993). A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game. *Nature*, *364*(6432), 56–58. https://doi.org/10.1038/364056a0

Nowak, M. A., & Sigmund, K. (2005). Evolution of indirect reciprocity. *Nature*, *437*(7063), 1291–1298. https://doi.org/10.1038/nature04131

Ochs, J. (1990). The coordination problem in decentralized markets: An experiment. *The Quarterly Journal of Economics*, *105*(2), 545–559. https://doi.org/10.2307/2937800

Ostrom, E. (1998). A behavioral approach to the rational choice theory of collective action: Presidential address, American Political Science Association, 1997. *American Political Science Review*, *92*(1), 1–22. https://doi.org/10.2307/2585925

Paolacci, G., & Yalcin, G. (2020). Fewer but poorer: Benevolent partiality in prosocial preferences. *Judgment and Decision Making*, *15*(2), 173–181. http://journal.sjdm.org/19/191220a/jdm191220a.pdf

Parker, M. T., & Janoff-Bulman, R. (2013). Lessons from morality-based social identity: The power of outgroup "hate,"not just ingroup "love". *Social Justice Research*, *26*(1), 81–96. https://doi.org/10.1007/s11211-012-0175-6

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1988). Adaptive strategy selection in decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *14*(3), 534–552. https://doi.org/10.1037/0278-7393.14.3.534

Penner, L. A., Dovidio, J. F., Piliavin, J. A., & Schroeder, D. A. (2005). Prosocial behavior: Multilevel perspectives. *Annual Review of Psychology*, *56*, 365–392. https://doi.org/10.1146/annurev.psych.56.091103.070141

Persad, G., Wertheimer, A., & Emanuel, E. J. (2009). Principles for allocation of scarce medical interventions. *The Lancet*, *373*(9661), 423–431. https://doi.org/10.1016/S0140-6736(09)60137-9

Peterson, C. R., & Ulelha, Z. J. (1965). Sequential patterns and maximizing. *Joural of Experimental Psychology*, *69*(1), 1–4. https://doi.org/10.1037/h0021597

Piazza, J., & Sousa, P. (2014). Religiosity, political orientation, and consequentialist moral thinking. *Social Psychological and Personality Science*, *5*(3), 334–342. https://doi.org/10.1177/1948550613492826

Puskaric, M., von Helversen, B., & Rieskamp, J. (2018). How social information affects information search and choice in probabilistic inferences. *Acta Psychologica*, *182*, 166–176. https://doi.org/10.1016/j.actpsy.2017.08.004

R Core Team. (2018). R: A Language and Environment for Statistical Computing. https://www.r-project.org/

Raafat, R. M., Chater, N., & Frith, C. (2009). Herding in humans. *Trends in Cognitive Sciences*, *13*(10), 420–428. https://doi.org/10.1016/j.tics.2009.08.002

Rader, C. A., Larrick, R. P., & Soll, J. B. (2017). Advice as a form of social influence: Informational motives and the consequences for accuracy. *Social and Personality Psychology Compass*, *11*(8), 1–17. https://doi.org/10.1111/spc3.12329

Rai, T. S., & Fiske, A. P. (2011). Moral psychology is relationship regulation: Moral motives for unity, hierarchy, equality, and proportionality. *Psychological Review*, *118*(1), 57–75. https://doi.org/10.1037/a0021867

Rapoport, A., Selae, D. A., & Ordóñez, L. (2002). Tacit coordination in choice between certain outcomes in endogenously determined lotteries. *Journal of Risk and Uncertainty*, *25*(1), 21–45. https://doi.org/10.1023/A:1016315329973

Rendell, L., Boyd, R., Cownden, D., Enquist, M., Eriksson, K., Feldman, M. W., Fogarty, L., Ghirlanda, S., Lillicrap, T., & Laland, K. N. (2010). Why copy others? Insights from the social learning strategies tournament. *Science*, *328*(5975), 208–213. https://doi.org/10.1126/science.1184719

Reverdy, P. B., Srivastava, V., & Leonard, N. E. (2014). Modeling human decision making in generalized gaussian multiarmed bandits. *Proceedings of the IEEE*, *102*(4), 544–571. https://doi.org/10.1109/JPROC.2014.2307024

Ried, K., Müller, T., & Briegel, H. J. (2019). Modelling collective motion based on the principle of agency: General framework and the case of marching locusts. *PLoS ONE*, *14*(2), 1–21. https://doi.org/10.1371/journal.pone.0212044

Rieucau, G., & Giraldeau, L.-A. (2011). Exploring the costs and benefits of social information use: An appraisal of current experimental evidence. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *366*(1567), 949–957. https://doi.org/10.1098/rstb.2010.0325

Robbins, H. (1952). Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, *58*(5), 527–536. https://doi.org/10.1090/S0002-9904-1952-09620-8

Roberts, M. E., & Goldstone, R. L. (2011). Adaptive group coordination and role differentiation. *PLoS ONE*, *6*(7), e22377. https://doi.org/10.1371/journal.pone.0022377

Rogers, A. R. (1988). Does Biology Constrain Culture? *American Anthropologist*, *90*(4), 819–831. http://www.jstor.org/stable/680759

Salganik, M. J., Dodds, P. S., & Watts, D. J. (2006). Experimental study of inequality and unpredictability in an artificial cultural market. *Science*, *311*(5762), 854–856. https://doi.org/10.1126/science.1121066

Samuelson, W., & Zeckhauser, R. (1988). Status quo bias in decision making. *Journal of Risk and Uncertainty*, *1*(1), 7–59. https://doi.org/10.1007/BF00055564

Sasaki, T., & Pratt, S. C. (2011). Emergence of group rationality from irrational individuals. *Behavioral Ecology*, *22*(2), 276–281. https://doi.org/10.1093/beheco/arq198

Savage, L. J. (1954). *The foundations of statistics*. Wiley.

Savulescu, J., Cameron, J., & Wilkinson, D. (2020). Equality or utility? Ethics and law of rationing ventilators. *British Journal of Anaesthesia*, *125*. https://doi.org/10.1016/j.bja.2020.04.011

Schaich Borg, J., Hynes, C., Van Horn, J., Grafton, S., & Sinnott-Armstrong, W. (2006). Consequences, action, and intention as factors in moral judgments: An fMRI investigation. *Journal of Cognitive Neuroscience*, *18*(5), 803–817. https://doi.org/10.1162/jocn.2006.18.5.803

Schrah, G., Dalal, R., & Sniezek, J. (2006). The adaptive decision-maker is not an island: Integrative expert advice with information search. *Journal of Behavioral Decision-Making*, *19*, 43–60. https://doi.org/10.1002/bdm.514

Schultze, T., Mojzisch, A., & Schulz-Hardt, S. (2017). On the inability to ignore useless advice. *Experimental Psychology*, *64*(3), 170–183. https://doi.org/10.1027/1618-3169/a00036

Schulz, E., Konstantinidis, E., & Speekenbrink, M. (2018). Putting bandits into context: How function learning supports decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *44*(6), 927–943. https://doi.org/10.1037/xlm0000463

Schulz, E., Wu, C. M., Ruggeri, A., & Meder, B. (2019). Searching for rewards like a child means less generalization and more directed exploration. *Psychological Science*, *30*(11), 1561–1572. https://doi.org/10.1177/0956797619863663

Schulze, C., Gaissmaier, W., & Newell, B. R. (2020). Maximizing as satisficing: On pattern matching and probability maximizing in groups and individuals. *Cognition*, *205*, 104382. https://doi.org/10.1016/j.cognition.2020.104382

Schulze, C., & Newell, B. R. (2015). Compete, coordinate, and cooperate: How to exploit uncertain environments with social interaction. *Journal of Experimental Psychology: General*, *144*(5), 967–981. https://doi.org/10.1037/xge0000096

Schulze, C., & Newell, B. R. (2016). More heads choose better than one: Group decision making can eliminate probability matching. *Psychonomic Bulletin & Review*, *23*(3), 907–914. https://doi.org/10.3758/s13423-015-0949-6

Schulze, C., van Ravenzwaaij, D., & Newell, B. R. (2015). Of matchers and maximizers: How competition shapes choice under risk and uncertainty. *Cognitive Psychology*, *78*, 78–98. https://doi.org/10.1016/j.cogpsych.2015.03.002

Seeley, T. D. (2003). Consensus building during nest-site selection in honey bee swarms: The expiration of dissent. *Behavioral Ecology and Sociobiology*, *53*(6), 417–424. https://doi.org/10.1007/s00265-003-0598-z

Selten, R., Chmura, T., Pitz, T., Kube, S., & Schreckenberg, M. (2007). Commuters route choice behaviour. *Games and Economic Behavior*, *58*(2), 394–406. https://doi.org/10.1016/j.geb.2006.03.012

Sherry, D. F., & Galef, B. G. (1984). Cultural transmission without imitation: Milk bottle opening by birds. *Animal Behaviour*, *32*(3), 937–938. https://doi.org/10.1016/S0003-3472(84)80185-2

Shirado, H., & Christakis, N. A. (2017). Locally noisy autonomous agents improve global human coordination in network experiments. *Nature*, *545*(7654), 370–374. https://doi.org/10.1038/nature22332

Sidowski, J. B. (1957). Reward and punishment in a minimal social situation. *Journal of Experimental Psychology*, *54*(5), 318–326. https://doi.org/10.1037/h0047119

Sidowski, J. B., Wyckoff, L. B., & Tabory, L. (1956). The influence of reinforcement and punishment in a minimal social situation. *The Journal of Abnormal and Social Psychology*, *52*(1), 115–119. https://doi.org/10.1037/h0045737

Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, *63*(2), 129–138. https://doi.org/10.1037/h0042769

Simonsohn, U. (2009). Direct risk aversion: Evidence from risky prospects valued below their worst outcome. *Psychological Science*, *20*(6), 686–692. https://doi.org/10.1111/j.1467-9280.2009.02349.x

Sims, C. R., Neth, H., Jacobs, R. A., & Gray, W. D. (2013). Melioration as rational choice: Sequential decision making in uncertain environments. *Psychological Review*, *120*(1), 139–154. https://doi.org/10.1037/a0030850

Singer, P. (2005). Ethics and intuitions. *The Journal of Ethics*, *9*(3-4), 331–352. https://doi.org/10.1007/s10892-005-3508-y

Singer, P. (2011). *The expanding circle: Ethics, evolution, and moral progress*. Princeton University Press.

Sniezek, J. A., & Buckley, T. (1995). Cueing and cognitive conflict in judge-advisor decision making. *Organizational Behavior and Human Decision Processes*, *62*(2), 159–174. https://doi.org/10.1006/obhd.1995.1040

Sokolowski, M. B. C., Tonneau, F., & Freixa i Baqué, E. (1999). The ideal free distribution in humans: An experimental test. *Psychonomic Bulletin & Review*, *6*(1), 157–161. https://doi.org/10.3758/BF03210824

Solnica, A., Barski, L., & Jotkowitz, A. (2020). Allocation of scarce resources during the COVID-19 pandemic: A jewish ethical perspective. *Journal of Medical Ethics*. https://doi.org/10.1136/medethics-2020-106242

Steyvers, M., Lee, M. D., & Wagenmakers, E.-J. (2009). A Bayesian analysis of human decision-making on bandit problems. *Journal of Mathematical Psychology*, *53*, 168–179. https://doi.org/10.1016/j.jmp.2008.11.002

Stijnen, M. M. N., & Dijker, A. J. M. (2011). Reciprocity and need in posthumous organ donation: The mediating role of moral emotions. *Social Psychological and Personality Science*, *2*(4), 387–394. https://doi.org/10.1177/1948550610393749

Sunstein, C. R. (2005). Moral heuristics. *Behavioral and Brain Sciences*, *28*(4), 531–541. https://doi.org/10.1017/S0140525X0534009X

Surowiecki, J. (2004). *The Wisdom of Crowds*. Random House.

Sutton, R. S., & Barto, A. (1998). *Reinforcement Learning: An Introduction*. The MIT Press.

Tajfel, H., Billig, M. G., Bundy, R. P., & Flament, C. (1971). Social categorization and intergroup behaviour. *European Journal of Social Psychology*, *1*(2), 149–178. https://doi.org/10.1002/ejsp.2420010202

Tajfel, H., & Turner, J. C. (1979). An integrative theory of intergroup conflict. In W. G. Austin & S. Worchel (Eds.), *The social psychology of intergroup relations* (pp. 33–47). Brooks Cole.

Tennie, C., Frith, U., & Frith, C. D. (2010). Reputation management in the age of the world-wide web. *Trends in Cognitive Sciences*, *14*(11), 482–488. https://doi.org/10.1016/j.tics.2010.07.003

Theraulaz, G., & Bonabeau, E. (1999). A brief history of stigmergy. *Artificial Life*, *5*(2), 97–116. https://doi.org/10.1162/106454699568700

Thomson, J. J. (1985). The trolley problem. *The Yale Law Journal*, *94*(6), 1395–1415. https://doi.org/10.2307/796133

Thuijsman, F., Peleg, B., Amitai, M., & Shmida, A. (1995). Automata, matching and foraging behavior of bees. *Journal of Theoretical Biology*, *175*(3), 305–316. https://doi.org/10.1006/jtbi.1995.0144

Tobler, P. N., Kalis, A., & Kalenscher, T. (2008). The role of moral utility in decision making: An interdisciplinary framework. *Cognitive, Affective, & Behavioral Neuroscience*, *8*(4), 390–401. https://doi.org/10.3758/CABN.8.4.390

Toplak, M. E., West, R. F., & Stanovich, K. E. (2014). Assessing miserly information processing: An expansion of the cognitive reflection test. *Thinking & Reasoning*, *20*(2), 147–168. https://doi.org/10.1080/13546783.2013.844729

Toyokawa, W., Kim, H. R., & Kameda, T. (2014). Human collective intelligence under dual exploration-exploitation dilemmas. *PLoS ONE*, *9*(4), 1–9. https://doi.org/10.1371/journal.pone.0095789

Toyokawa, W., Whalen, A., & Laland, K. N. (2019). Social learning strategies regulate the wisdom and madness of interactive crowds. *Nature Human Behaviour*, *3*(2), 183–193. https://doi.org/10.1038/s41562-018-0518-x

Uhlmann, E. L., Pizarro, D. A., Tannenbaum, D., & Ditto, P. H. (2009). The motivated use of moral principles. *Judgment and Decision Making*, *4*(6), 479–491. http://journal.sjdm.org/9616/jdm9616.pdf

Van Bavel, J. J., FeldmanHall, O., & Mende-Siedlecki, P. (2015). The neuroscience of moral cognition: From dual processes to dynamic systems. *Current Opinion in Psychology*, *6*, 167–172. https://doi.org/10.1016/j.copsyc.2015.08.009

van den Bos, R., Jolles, J. W., & Homberg, J. R. (2013). Social modulation of decision-making: A cross-species review. *Frontiers in Human Neuroscience*, *7*(June), 1–16. https://doi.org/10.3389/fnhum.2013.00301

Volz, K. G., & Gigerenzer, G. (2012). Cognitive processes in decisions under risk are not the same as in decisions under uncertainty. *Frontiers in Neuroscience*, *6*(105), 1–6. https://doi.org/10.3389/fnins.2012.00105

von Dawans, B., Fischbacher, U., Kirschbaum, C., Fehr, E., & Heinrichs, M. (2012). The social dimension of stress reactivity: Acute stress increases prosocial behavior in humans. *Psychological Science*, *23*(6), 651–660. https://doi.org/10.1177/0956797611431576

Vulkan, N. (2000). An economist's perspective on probability matching. *Journal of Economic Surveys*, *14*(1), 101–118. https://doi.org/10.1111/1467-6419.00106

Waldmann, M. R., Nagel, J., & Wiegmann, A. (2012). Moral judgment. In K. J. Holyoak & R. G. Morrison (Eds.), *Oxford library of psychology. The Oxford handbook of thinking and reasoning* (pp. 364–389). Oxford University Press.

West, R. F., & Stanovich, K. E. (2003). Is probability matching smart? Associations between probabilistic choices and cognitive ability. *Memory & Cognition*, *31*(2), 243–251. https://doi.org/10.3758/BF03194383

White, D. B., & Lo, B. (2020). A framework for rationing ventilators and critical care beds during the COVID-19 pandemic. *JAMA*, *323*(18), 1773–1774. https://doi.org/10.1001/jama.2020.5046

Wilke, A., Scheibehenne, B., Gaissmaier, W., McCanney, P., & Barrett, H. C. (2014). Illusionary pattern detection in habitual gamblers. *Evolution and Human Behavior*, *35*(4), 291–297. https://doi.org/10.1016/j.evolhumbehav.2014.02.010

Wilkinson, D., Zohny, H., Kappes, A., Sinnott-Armstrong, W., & Savulescu, J. (2020). Which factors should be included in triage? An online survey of the attitudes of the UK general public to pandemic triage dilemmas. *BMJ Open*, *10*(12). https://doi.org/10.1136/bmjopen-2020-045593

Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A., & Cohen, J. D. (2014). Humans use directed and random exploration to solve the explore–exploit dilemma. *Journal of Experimental Psychology: General*, *143*(6), 2074–2081. https://doi.org/10.1037/a0038199

Wolford, G., Newman, S. E., Miller, M. B., & Wig, G. S. (2004). Searching for patterns in random sequences. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale*, *58*(4), 221–228. https://doi.org/10.1037/h0087446

Wood, L. A., Kendal, R. L., & Flynn, E. G. (2013). Copy me or copy you? The effect of prior experience on social learning. *Cognition*, *127*(2), 203–213. https://doi.org/10.1016/j.cognition.2013.01.002

Wu, C. M., Schulz, E., Speekenbrink, M., Nelson, J. D., & Meder, B. (2018). Generalization guides human exploration in vast decision spaces. *Nature Human Behaviour*, *2*(12), 915–924. https://doi.org/10.1038/s41562-018-0467-4

Yahosseini, K. S., & Moussaïd, M. (2020). Comparing groups of independent solvers and transmission chains as methods for collective problem-solving. *Scientific Reports*, *10*(1), 3060. https://doi.org/10.1038/s41598-020-59946-9

Yamagishi, T., & Kiyonari, T. (2000). The group as the container of generalized reciprocity. *Social Psychology Quarterly*, *63*(2), 116–132. https://doi.org/10.2307/2695887

Yaniv, I. (2004). Receiving other people's advice: Influence and benefit. *Organizational Behavior and Human Decision Processes*, *93*(1), 1–13. https://doi.org/10.1016/j.obhdp.2003.08.002

Yaniv, I., & Kleinberger, E. (2000). Advice taking in decision making: Egocentric discounting and reputation formation. *Organizational Behavior and Human Decision Processes*, *83*(2), 260–281. https://doi.org/10.1006/obhd.2000.2909

Yaniv, I., & Milyavsky, M. (2007). Using advice from multiple sources to revise and improve judgments. *Organizational Behavior and Human Decision Processes*, *103*(1), 104–120. https://doi.org/10.1016/j.obhdp.2006.05.006