Cognitive processes in social and moral decision making

Rahal, R.M.

Chapter 1

Introduction

Every day, we face uncountable big and small decisions. From deciding whether to surprise your spouse with breakfast in bed or sleeping in yourself, deciding between giving honest feedback during a shopping spree or sparing your friend's feelings, deciding whether to help a stranger at the post office pick up the mail they dropped, to deciding to let someone go in front of you in the queue in the supermarket because they only have two items in their basket, navigating the normal course of life means deciding what to do time and time again. Given their omnipresence and importance for shaping our everyday experiences, the decisions we make have roused the interest of research at the crossroads of psychology, economics, philosophy and beyond.

In this dissertation, the focus is on decisions set in two specific contexts: (i) decisions in social dilemmas in group contexts, and (ii) decisions in moral dilemmas. In the former, decision makers choose whether to be prosocial or selfish towards in- or outgroup members, while in the latter, decision makers choose between what they think is morally right or wrong. Gaining deeper understanding of the decisions made in both contexts under investigation here has immediate pertinence for society. On the one hand, in a time of globalized mobility, strong emphasis on freedom of movement and of rising flows of migration societies experience new challenges in deciding how to face increasing numbers of immigrants of all types. For instance, as the number of newcomers rises, the native population's willingness to trust and help a foreigner might sink (La Macchia, Louis, Hornsey, & Leonardelli, 2016). The question how people interact with ingroup versus outgroup members and strangers, and potentially discriminate against newcomers, thus has gained increasing importance in

many societies. Although the question of how societies face increased immigration is a complex issue, parts of this challenge can be distilled into a question of intergroup interactions: how do members of an ingroup interact with those of an outgroup? What are the driving forces of cooperation in these circumstances? Understanding decisions embedded in a group setting therefore affords opportunities to encourage fairness and cooperation beyond the borders of one's ingroup.

On the other hand, the more abstract question asking what is morally right or wrong has occupied philosophers and laypeople alike for centuries. A variety of moral questions and the way they are decided determine the way we live together in society. For instance, is it morally right for doctors not reveal their diagnoses to terminally ill patients, in an effort to uphold the Hippocratic Oath to do no harm (Oken, 1961; Sullivan, Menapace, & White, 2001)? And is it morally right for a legal system to hinder people from committing voluntarily chosen behavior with limited externalities such as prostitution (Euchner & Knill, 2015), the consumption of soft-drugs (Blumenson & Nilsen, 2010) or suicide (Tännsjö, 2005)? Embedded in such broader moral questions, two distinct moral schools of thought are focused on in this dissertation: utilitarianism, the orientation to maximize outcomes in a moral struggle, and deontology, the orientation to uphold absolute rules. In a recent example of how these motivations are directly juxtaposed in moral decisions of societal scope, consider that in the aftermath of the 9/11 terror attacks the German parliament passed changes to the air security law, legalizing the shooting down of hijacked, weaponized airplanes (Gesetz zur Neuregelung von Luftsicherheitsaufgaben, 2005). Therein, the parliament subscribed to a utilitarian perspective of sacrificing a small number of kidnapped people to save lives on the ground. The German constitutional court, however, overturned the bill (Bundesverfassungsgericht, 2006), citing that quantifying human lives violates the guarantee of human dignity of the German Basic Law (Grundgesetz, 1949). Therein, the judicial review reminded the society of its basic values, and in this case of its core deontological values. Here, rather than aiming to

understand complex networks of morality and moral reasoning on the societal scale, the question is again distilled to individual decision makers as the smallest observable unit in these situations. How do people decide which solution to moral dilemmas like the air security law they support? Is moral reasoning the driving force in these circumstances or do people rely on their moral intuitions to solve these problems?

In both areas, a large body of research in psychology, behavioral economics and beyond, which is reviewed below, has been devoted to understanding the situational and person-specific influences on choice outcomes. For instance, who decides to help others and who does not, and under which circumstances are people more likely to make utilitarian decisions? Surpassing this strong focus on choice outcomes, current debates about prosocial and moral decision making, however, have taken a turn towards the cognitive processes driving choices. This literature, which is summarized below, addresses the question how certain choices are made: Are some choices easier to make than others? Do some decision makers inform their choices better, and which information is more important to different decision makers? Answering such questions not only posits the opportunity for gaining a better understanding of the mechanisms through which people arrive at their decisions, but also offers possibilities to design informed interventions promoting fair and morally good decisions. Therefore, in line with current developments in scientific interest, this dissertation investigates in the cognitive processes involved in prosociality and morality.

Gaps in the literature about these decision processes are identified, and subsequently addressed in the empirical chapters of this dissertation. Specifically, regarding prosociality in intergroup settings, three questions are addressed: First, what determines whether decision makers acquire information about others' group membership when deciding whether to make prosocial or selfish choices when facing in- vs. outgroup members? Second, differences in utility derived from in- vs. outgroup

members' outcomes reflected in differences in attention towards own and others' outcomes? And third, are there differences in how much effort decision makers invest in informing their decision to be prosocial or selfish when facing in- vs. outgroup members? Regarding moral decision making, we ask whether deontological vs. utilitarian decision makers differ in the way they make moral decisions, and whether these differences are in line with a dual-process theory of moral decision making, or with a choice discriminability account.

As tools to facilitate the investigation of decision processes, process tracing techniques, and in particular eye tracking, are introduced and critically discussed. Moreover, process theories of decision making are introduced, to provide context for the subsequent application of process predictions from the attentional drift diffusion model to the research questions outlined above.

Finally, implications of the present research are discussed. In the domain of prosociality, such insights offer unique opportunities for tailored interventions to promote fairness. In the domain of morality, insights about processing during decision making yield immediate implications for current theoretical debates. Lastly, this dissertation discusses limitations of the studies presented, and summarizes avenues for future research.

Decision Making

The research presented in this dissertation is embedded in the broader decision making literature. Research on judgment and decision making focuses on "investigat[ing] the processes by which people draw conclusion, reach evaluations, and make choices." (Goldstein & Hogarth, 1997). To give an overview of the field, a brief historical overview must begin at the roots of modern investigations of individual choice: Expected Utility Theory (von Neumann & Morgenstern, 1944) and Subjective Expected Utility Theory (Savage, 1954). Given the assumptions of well-defined,

transitive, continuous and independent preferences, the theory proposed that decision makers would behave as if they maximized expected utility over possible outcomes. Decision makers' utility was conceptualized as related to decision makers' desire for certain outcomes, and argued to be captured by observing or predicting choices (Samuelson, 1938). Soon, this conceptualization of decision making spread from economics to other disciplines such as psychology, spurred forward by publications such as Edwards' (1954) seminal paper and Luce and Raiffa's (1957) book introducing game and decision theory. Psychological research responded to Expected Utility Theory by questioning its descriptive validity and, consequently, its theoretical preeminence, drawing on paradoxes (e.g., Allais, 1953; Ellsberg, 1961; for a historical overview, see MacCrimmon & Larsson, 1979) and experiments showing the violation of its predictions and of its axioms (for an overview, see Goldstein & Hogarth, 1997; for historic reviews, see Amnon Rapoport & Wallsten, 1972; Slovic, Fischhoff, & Lichtenstein, 1977).

This critical attitude also addressed the fundamental assumption of Expected Utility Theory that actors would behave rationally to maximize their own utility. Expected utility theory leads to the prediction that decision makers choose what is best for them, regardless of the outcomes of others. For instance, Selten (1978) had theoretically demonstrated via backward induction that defection is the subgame perfect equilibrium in a two-player repeated prisoner's dilemma game. Yet, empirical evidence showed that decision makers do not defect fully (e.g., Anatol Rapoport & Dale, 1966; Selten & Stoecker, 1986; for contemporary overviews, see van Lange, Joireman, Parks, & Van Dijk, 2013; Chaudhuri, 2011), violating the expectations derived from Expected Utility Theory.

Social Preferences and Prosociality

Breaking with the assumption that decision makers will always do what maximizes their own material outcomes, a new strand of the literature evolved, which

considered that others' outcomes also matter to decision makers. This literature on social preferences is the basis for scientific debate about prosociality, and therefore briefly reviewed here. By introducing the concept of social preferences, others' outcomes were integrated into decision makers' utility functions (e.g., Axelrod & Hamilton, 1981; Dawes, 1980; Taylor, 1977). In other words, this line of research began investigating the nature and boundary conditions of prosocial acts, i.e., actions intended to benefit others motivated by the concern for their welfare (e.g., Mussen & Eisenberg-Berg, 1977) even at the cost of decision makers' own material outcomes.

In social decision settings capturing situations where individuals' choices affect not only themselves but also others, formal models of Utility Theory are used to describe how individuals value certain actions depending on their weighting of outcomes for themselves and others affected (Messick & McClintock, 1968). An example of such a decision situation is a dictator game (Forsythe, Horowitz, Savin, & Sefton, 1994), where one person decides about both her own outcomes and those of a second player. For simplicity and in line with the materials used in Chapters 1 and 2, a decomposed dictator game is introduced here. Decision makers are assigned to the role of the dictator or the receiver. The dictator chooses between two options, where one alternative maximizes her own monetary payoff (for example, in Option A, the dictator receives 10€ and the matched player receives 2€) and the other option benefits the other player through a reduction of the dictator's payoff (for example, in Option B, the dictator receives 7€ and the matched player receives 5€). The tradeoff between own and others' payoff can be represented as differences in the respective decision weights w_{own} and w_{other}:

$$U = w_{own} \times (\text{own payoff}) + w_{other} \times (\text{others' payoff}).^{1} \quad (1)$$

Individuals who value their own payoffs highly and disregard others' outcomes (e.g., w_{own} = 1; w_{other} = 0) are more likely to choose the option maximizing their own payoff, which has higher a utility for them ($Payoff_{Option\ A}$ = 10, $Payoff_{Option\ B}$ = 7) compared to individuals who assign high value to both their own and others' payoffs (e.g., w_{own} = 0.5; w_{other} = 0.5, $Payoff_{Option\ A}$ = 6, $Payoff_{Option\ B}$ = 6.5). This preference for allocating resources between oneself and another person is thought to have relatively stable, trait-like characteristics (e.g., McClintock & van Avermaet, 1982; Swap & Rubin, 1983; van Lange, 2000), referred to as Social Value Orientation (SVO, van Lange, 1999, for an overview, see Murphy & Ackermann, 2013). The weight assigned to own and others' outcomes allows the categorization of decision makers into types. Most commonly, decision makers present with altruistic (w_{own} = 0; w_{other} = 1), prosocial (w_{own} = 0.5; w_{other} = 0.5), individualistic (w_{own} = 1; w_{other} = 0) or competitive (w_{own} = 0.5; w_{other} = -0.5) preferences, or intermediate assigned weights (Murphy, Ackermann, & Handgraaf, 2011), although other weightings of own and others' outcomes are possible (e.g., martyrdom: w_{own} = -0.5; w_{other} = 0.5). SVO has been shown to be strongly related to the Honesty-Humility factor of personality (Hilbig & Zettler, 2009), the general tendency to be fair to others (Ashton & Lee, 2007), and strongly predicts cooperation behavior in a variety of situations (for a review, see Bogaert, Boone, & Declerck, 2008): Prosocial and altruistic types are more likely to share resources than individualistic and competitive types.

Prosocial behavior, however, is not only determined by inter-individual differences in the degree to which others' outcomes are valued, but also driven by

[1] The form in which others' outcomes are integrated in to the decision makers' utility function has been the object of several theories of social preferences. Depending on the situations where behavior is predicted, these theories added different components to the simple utility function suggested above (see Bolton & Ockenfels, 2000; Fehr & Schmidt, 1999; Rabin, 1993).

situation-specific influences on the weights assigned to others' outcomes (for an overview, see Kopelman, Weber, & Messick, 2002). Prominently featured among such situational influences on prosocial behavior are group settings.

Prosociality in Group Settings

As Chapters 2 and 3 are focused on prosociality in group settings, the theoretical background on this topic is reviewed below, including perspectives from both economics and psychology. Phenomenologically, the influence of group membership on prosocial behavior has been observed in a variety of contexts, indicating that ingroup members are often favored (Hewstone, Rubin, & Willis, 2002). For instance, ingroup members are trusted more (Glaeser, Laibson, Scheinkman, & Soutter, 2000), ingroup cooperation is rewarded with higher reciprocity (Chen & Li, 2009), and ingroup members are favored in prosocial decision making (for a review, see Balliet, Wu, & De Dreu, 2014).

A model of how prosociality is affected by group membership was proposed by Chen and Li (2009), where the weight w_{other} assigned to others' outcomes in decision makers' utility functions is determined by their membership to the in- our outgroup *I*, which leads to the (simplified) utility function

$$U = w_{own} \times (\text{own payoff}) + w^{I}_{other} \times (\text{others' payoff}). \quad (2)$$

In more detail (Chen & Li, 2009, p. 441), the utility function proposes that the weight placed on others' outcomes depending on their group membership is differently affected when decision makers' own outcomes are larger or smaller than the other persons' outcomes. Supporting this theoretical proposition empirically, decision makers were found to be more charitable to in- vs. outgroup members when their own outcomes were larger than the other players', and to be more envious of out- vs. ingroup members when their own outcomes were smaller than the other players' (Chen & Li, 2009).

Beyond this conceptualization of how group membership influences prosociality, a number of prominent theories are concerned with explaining human interactions in group settings more broadly (for an overview, see Hogg, 2016). In contemporary research, three theories are predominantly drawn from to understand intergroup behavior, which are summarized below.

Social Identity Theory (Tajfel & Turner, 1979, 1986) summarizes cognitive, motivational and behavioral influences of group settings. Aiming to explain why people form and maintain groups, the Social Identity Theory posited that individuals derive part of their self-concept from their categorization into social groups. Therein, they act not as individuals but as (self-categorized) members of their groups who aim to achieve positive distinctiveness from other groups (Tajfel & Turner, 1979). To maintain the ingroup identity and distinctiveness from other groups, discriminatory behavior emerges. It was argued that behavior favoring the ingroup and discriminating against the outgroup does not require competition of groups over resources (in contrast to the Realistic Group Conflict Theory), but that mere categorization into groups is a sufficient condition. When decision makers were assigned to minimal groups, i.e., groups with whom the participant shared no common past or future, group members were anonymous, and there was no conflict of interest over group outcomes, their choices demonstrated ingroup favoritism (e.g., Tajfel, 1970; Tajfel, Billig, Bundy, & Flament, 1971).

Translating the concepts of social categorization and identity into economic analyses, ingroup favoritism is also captured by Identity Theory (Akerlof & Kranton, 2000; Bénabou & Tirole, 2011). The theory holds that social categorization and the resulting identity, such as being an alumnus of a certain college, is associated with specific expectations of behavior, e.g., for donations to be directed to one's alma mater, not a different college (Akerlof & Kranton, 2000). Complying with these expectations maximizes utility, while deviating from the expectations causes disutility.

Therefore, rational decision makers would behave in line with expectations stemming from the social categories they belong to.

While these theories assigned great importance to social identity in determining behavior in broader contexts, the Bounded Generalized Reciprocity Theory (Yamagishi, Jin, & Kiyonari, 1999) focused on the strategic advantage of group membership for maximizing one's own outcomes. The theory proposed that groups help decision makers form better expectations of others' likelihood to cooperate. Ingroup members can be expected to cooperate more than outgroup members, in turn increasing the likelihood that other ingroup members will cooperate with them. These expectations are based on a system of indirect reciprocity and reputation. The theory proposes that ingroup members have a reputation to cooperate that leads other ingroup members to expect them to cooperate. Additionally, reputation promotes cooperation even when there is no direct history of cooperativeness between the same people. Decision makers obtain and maintain the ingroup member status that allows them to benefit from favorable treatment by the ingroup. Critically, ingroup favoritism depends on the knowledge of others' group membership, and the expectation that ingroup members cooperate more, as demonstrated in Yamagishi and Kiyonari (2000).

In sum, in several theoretical systems and in many demonstrations of behavioral manifestations, decision makers favor the ingroup, in particular by being more prosocial to ingroup members than to outgroup members. Therein, ingroup favoritism constitutes a special case of prosocial behavior.

Cognitive Processes driving Prosociality

Despite the originally behavior-focused approach to decision making inspired by rational choice theories in economics (Gul & Pesendorfer, 2008), a more recent literature has turned towards investigating how decisions are reached (Camerer,

Loewenstein, & Prelec, 2004; for an overview, see Krajbich & Dean, 2015). Mirroring this growing interested in the cognitive processes of decision making generally, in the area of prosociality, advances have been made to look inside the black box of social preferences. A short overview of this literature (see below) reveals gaps in our understanding in particular when group settings are concerned.

In individual-level behavior, the current debate on underlying cognitive processes in prosociality draws on the more general dual-process view of decision making, distinguishing fast and intuitive processing from slow and deliberate decision strategies (Gigerenzer & Selten, 2001; Kahneman, 2003). On the one hand, some argue that cooperation is intuitive (Social Heuristics Hypothesis; Rand et al., 2014). This notion is based on findings that higher contributions in a Public Goods Game were made more quickly, and that time pressure manipulations increased contributions in comparison to time delay (Rand, Greene, & Nowak, 2012).[2] However, diverging findings showed that cooperation was only effortless for prosocial decision makers, whose preference is consistent with the choice outcome (Mischkowski & Glöckner, 2016). Moreover, when accounting for choice discriminability, reaction time differences between prosocial and selfish choices would no longer remain significant (Krajbich, Bartling, Hare, & Fehr, 2015). When manipulating choice difficulty, time pressure no longer consistently increased prosociality (Merkel & Lohse, 2018). Finally, extreme choices were found to be more effortless than intermediate choices where utility differences are smaller (A. M. Evans, Dillon, & Rand, 2015). Therefore, in parallel to the debate in moral decision making, a competing choice discriminability account

[2] Although the finding could be reproduced in several contexts (for a meta-analytic overview, see Rand, 2016), some diverging evidence showed null effects (Tinghög et al., 2013; Verkoeijen & Bouwmeester, 2014). Moreover, other evidence suggested that defection was the more intuitive choice. For instance, prosocial SVO was associated with less intuitive processing through longer response times, higher fixation counts, and a higher proportion of inspected pieces of information (Fiedler, Glöckner, Nicklisch, & Dickert, 2013), a pattern that is replicated in Chapter 3 of this dissertation.

has emerged in prosocial decision making, proposing that decisions where preferences allow a clear distinction between the alternatives are processes more effortlessly.

Although the debate about cognitive processes involved in prosocial choices in general is still ongoing, a socio-cognitive strand of the literature has added first empirical investigations of the cognitive bases specific to prosociality in group contexts (for a review, see Amodio, 2014). Initial evidence suggests that mental configurations for processing information related to in- and outgroup members differ. For example this evidence shows that ingroup members are evaluated more positively (Brewer, 1979; LeVine & Campbell, 1972; Mullen, Brown, & Smith, 1992; Perdue, Dovidio, Gurtman, & Tyler, 1990), that people make positive spontaneous trait inferences about ingroup members (Otten & Moskowitz, 2000), and that positive stereotypes are more rapidly associated with ingroup members (Dovidio, Evans, & Tyler, 1986). Furthermore, ingroup members are judged to be more similar to oneself than outgroup members even in artificial, minimal groups (Cadinu & Rothbart, 1996), which is thought to be based on inferences from own traits to the ingroup (Otten & Epstude, 2006). In short, research finds that our cognitive landscape is set up in a way that facilitates ingroup favoring behavior.

In general, these research trajectories point to a growing research interest in cognitive processes in decision making with regard to prosociality generally and intergroup settings specifically. Issues such as what motivates decision makers to choose being selfish rather than prosocial, how much effort they invest into making their choice, and how difficult it is for them to arrive at their choice increasingly receive attention. However, the cognitive processes underlying prosociality in group settings are far from well described and well understood.

Investigating Cognitive Processes in Intergroup Settings

To advance the understanding of the cognitive processes involved in prosociality towards in- vs. outgroup members, two empirical chapters will be presented. In Chapter 2, we investigate knowledge of others' group membership as a necessary condition to engage in ingroup bias: Without knowing if the other person is an in- or outgroup member, decision makers cannot bias their choice depending on group membership. Finding out who belongs to one's own group or an outgroup, however, is not a given in everyday interactions. Rather, decision makers may sometimes not notice others' group membership, and sometimes even avoid this information. We show empirically that information about others' group membership is deliberately avoided by some decision makers, who stay blind to avoid biasing their choices. Evidence is presented both for ignorance to group membership information in visual information search via eye tracking, and for active ignorance in explicitly stated preferences. Moreover, we aimed to investigate if some individual characteristics would make certain decision makers more likely to remain ignorant to group membership information. We present evidence that high identification with the ingroup increased decision makers' likelihood to access group membership information, and that their first and second order beliefs of discrimination were correlated with ignorance. These insights suggest not only that systematic variation in whether people enter into a group-based decision problem or remain on an inter-individual interaction level has previously often been ignored, but also that custom-tailored interventions on discrimination beliefs may offer opportunities to ameliorate ingroup favoritism and outgroup derogation.

In Chapter 3, we investigate how decision makers weigh and process information to make decisions in a group setting. We show via eye tracking which importance decision makers with different social preferences assign to others' outcomes depending on their group membership, leading to ingroup favoritism and

discrimination against the outgroup. Moreover, we investigate if decision makers are biased to invest more effort to inform their decision when ingroup members were concerned. Evidence surprisingly but consistently over two eye tracking studies shows that only individualistic decision makers' decisions were relatively less effortful when making decisions affecting an outgroup member, although ingroup bias was small in this subsample. This finding suggests that individualists may struggle with competing group norms and individual preferences during the decision process, while prosocials more readily follow the group norm to bias their choices.

In sum, these chapters address the open question how decisions to be prosocial to in- vs. outgroup members are made. Therein, the chapters relate to an increasing interest in the cognitive mechanisms underlying decision making, and apply this perspective to intergroup settings.

Relation of Prosociality and Morality

The scientific trend towards uncovering underlying mechanisms of choice behavior is apparent not only in the area of prosociality, but also in the adjacent literature on morality. This similarity in the development of scientific investigations may be understood as a function of a time trend, but the strong relatedness of both research areas, which is reviewed below, may also play a role in determining research trajectories. As a theoretical and empirical concept, prosociality, be in interpersonal or intergroup situations, can be seen as embedded in the larger domain of morality. Moral decisions, in general, are decisions that are concerned with doing what is right or wrong (VandenBos, 2007), therein constituting a broader field of investigation in which prosociality is but one facet. Indeed, both theory and empirical practice have often considered prosociality as a facet of morality (for an in-depth discussion, consult Turiel, 2015). Regarding the theoretical perspective, evidence for the close relationship of prosociality and morality can be found in a variety of domains. For instance, in economic theories, morality is often alluded to in the shape of equitable

outcomes and fairness considerations (Bolton & Ockenfels, 2000; Fehr & Schmidt, 1999; Walster, Berscheid, & Walster, 1973). Moreover, theories of moral philosophy have often been concerned with human benevolence as a moral virtue (for overviews, see Kekes, 1987; Roberts, 1973). For example, Kant's (1797) thought experiment posed a dilemma between honesty and prosociality: a passerby must decide whether to reveal to a murderer the location the intended victim or to lie about it. Empirical investigations further support the notion that prosociality and morality are closely tied, by investigating prosociality in opposition to other moral virtues such as obedience of the law or honesty. For instance, weighing prosociality against honesty as in Kant's murderer dilemma is also evident in real-life decision situations: In a survey study, a sizeable proportion of physicians reported to lie to patients' insurance providers to help patients gain coverage for medical care (Wynia, Cummins, VanGeest, & Wilson, 2000). Further, Kohlberg (1981) investigated moral reasoning using the Heinz dilemma, in which a husband obtains a drug that could potentially save his wife's life by stealing it from a greedy druggist, such that decision makers must weight prosociality against property rights. Finally, research shows that people consider prosociality morally obligatory in contexts where competing selfish interests are small or absent (Kahn, 1992; Nucci & Turiel, 2009). Compellingly, a recent study also demonstrated that decision makers' preferences for doing what is moral, rather than specific preferences for equity and efficiency, determine prosociality (Capraro & Rand, 2017).

In sum, prosociality and morality can be understood as fundamentally interlinked concepts. Unsurprisingly, the investigation of moral decision making therefore raises similar research questions about the antecedents, influencing factors and processes of choices in the moral domain.

Morality

Moral decision making has been addressed from a variety of perspectives. A vast area of research is concerned with explaining behavior that is mostly agreed to be morally problematic, such as teaming up with another person to cheat the system so both of you will be financially better off (Weisel & Shalvi, 2015), or more morally ambiguous situations such as closing one's eyes to the consequences of one's decision on others to selfishly exploit moral wiggle room (Dana, Weber, & Kuang, 2007). Philosophical theories of morality explicitly ask how one can know what is right or wrong. Here, two general normative schools of thinking are of particular interest: deontology and utilitarianism. Deontological ethics is a normative theoretical position that judges the morality of an action based on rules, with strong ties to the Kantian argument of moral imperatives (Kant, 1785), while utilitarianism relies on the judgment of an action based on the utility (i.e., good, satisfaction, welfare, etc.) it produces (Hare, 1981; Mill, 1863). While the rules on which deontological judgments of actions are based can be seen as absolute (compare the literature on protected (Baron & Spranca, 1997) or sacred (Tetlock, Peterson, & Lerner, 1996) values, such as doing no harm, or being honest), utilitarianism requires weighing outcomes to achieve the greater good, suggesting its relation to Expected Utility Theory (Baron, 2014). Both perspectives prescribe which actions are right or wrong and can arrive at the same or at diverging evaluations of an action. Situations where deontology and utilitarianism arrive at opposing evaluations of the available choice options are referred to as moral dilemmas.

Empirical research aiming to juxtapose the deontology and utilitarianism as principles guiding decisions has made extensive use of these moral dilemmas. The classical example is the trolley dilemma (Foot, 1967), a hypothetical decision setting in which an out-of-control wagon is described to be hurling down train tracks, moments away from killing five unaware workers on the track. The decision maker is

a bystander and is presented with the opportunity to intervene in order to save the workers' lives by redirecting the wagon to another track. In doing so, however, one worker on the alternative track would be killed. Utilitarianism would prescribe sacrificing the one working, weighing the outcomes of this decision (one life lost) against those of the alternative (five lives lost). Deontology, however, would prescribe not rerouting the wagon, because taking lives is forbidden. A contemporary example of such a dilemma, the case of the German air safety law, was introduced above.

Empirical research shows that how people decide in moral dilemmas pitting deontological against utilitarian motives is influenced by a variety of factors. For instance, most such trolley-type moral dilemmas confound deontology with omission, through consistently linking the deontological option with non-action. As decision makers think that harm stemming from commission is worse than harm from omission (Spranca, Minsk, & Baron, 1991), they would rather choose an omitting act in line with the deontological perspective, and prefer harm to be done only as an indirect byproduct of a choice (Royzman & Baron, 2002). Moreover, when actions require personal force (Greene et al., 2009), decision makers are more hesitant to intervene by choosing the utilitarian option. In sum, there have been wide-ranging theoretical and empirical contributions in the field of moral decision making.

Cognitive Processes driving Moral Decisions

Surpassing the question what people will do in specific morally ambiguous situations, a substantive branch of the literature on morality has raised interest in the cognitive processes underlying moral decision making. Ethics has long been concerned with conflicts of moral theories and moral intuitions (e.g., Hare, 1981; Rawls, 1971): when our gut feeling tells us that an action is wrong, this assessment is not necessarily congruent with rational or theoretical deliberations. For example, many individuals believe that incest is morally wrong (Björklund, Haidt, & Murphy, 2000; Haidt, Koller, & Dias, 1993), even if it would be permissible from a utilitarian,

rationalized perspective when the behavior is a victimless non-procreative action between consenting adults and therefore has no negative consequences.

While some theories of morality posit that moral judgments take place solely on cognitive and rational bases, others emphasize the importance of moral intuitions. Exemplifying cognitive theories of morality, Kohlberg's (1963) theory was based on interviewing children of different ages, positing six stages of moral development summarized in three levels: pre-conventional morality, characterized by obedience to moral authorities, conventional morality, in which moral standards are internalized but not reflected, and post-conventional morality, where moral reasoning is based on individual cognitive judgments. In contrast, the Social Intuitionist Approach (Haidt, 2001) to moral judgment proposed that decision makers entirely rely on moral intuitions, which they justify and rationalize post hoc. Therein, the theory drew heavily on psychological literature on the boundaries of introspection (for an overview, see T. D. Wilson & Dunn, 2004), suggesting that people rationalize their behavior even when there is no objective possibility to rationalize it (e.g., Nisbett & Wilson, 1977; but see Smith & Miller, 1978 for a critique) or by providing reasons they had not actually used in their reasoning process (e.g., Meier, 1931).

Aiming to reconcile such opposing theories of moral judgment, Greene (2008) formulated a dual process model, connecting with the distinction of fast and intuitive processing from slow and deliberate decision strategies to more general theories of bounded rationality in decision making (for overviews, see Gigerenzer & Selten, 2001; Kahneman, 2003). The dual process model of moral judgment suggests that some moral decisions are supported by deliberate moral reasoning, while others are driven by emotional moral intuitions. Specifically, the theory posits the Central Tension Principle: utilitarian judgments, which require outcomes to be weighed to determine which one provides higher utility, require cognitive deliberation, while deontological judgments rely on automatic responses to moral issues (Greene, 2014).

In response to this dual process theory of moral decision making, a competing process model has been proposed (R. Kim et al., 2018), which may be able to better account for reaction time evidence used to infer the intuitive vs. deliberate processing mode. In this evidence accumulation model, choice discriminability is described as the central driver of the processing effort required to make a choice. The model proposes effortless responses in situations where the alternatives can be easily discriminated because a strong preference exists for one option. At the same time, more effortful responses are argued to stem from reduced choice discriminability given more similar preferences for the alternatives. In other words: when decision makers are undecided because they have no clear utilitarian or deontological preference, their responses in moral dilemmas would appear more deliberate, because more effort is required to make a choice.

Overall, the moral decision making literature has long harbored research interests about the underlying cognitive processes. Overcoming strictly behaviorist observations of what decision makers will choose, the literature has frequently addressed the question how decision makers choose a specific decision alternative.

Following this general research interest in cognitive processes underlying moral decision making, Chapter 4 of this dissertation contributes to the debate about intuitive vs. deliberate processing in deontological vs. utilitarian choices. We address the theoretical debate between the Dual Process Theory (Greene, 2007), suggesting that utilitarian decisions are supported by deliberate reasoning while deontological decisions are driven by moral intuitions, and the competing choice discriminability account (R. Kim et al., 2018), proposing that interactions of individual preferences and the choice made determine the processing effort needed. We use eye tracking to investigate both theoretical positions with relation to decision makers' attentional foci, their decision effort and conflictedness during the decision process. We develop third-party dictator game with a non-hypothetical moral good, in which choices and gaze

behavior are studied. Surprisingly, results indicated that to deontologists optimized outcomes were more important during the decision process than the violation rules, which one would rather expect from utilitarian decision makers. Moreover, the study showed that decision effort was lowest for decision makers with strong utilitarian moral preferences, while decision effort was higher among deontologists, in contrast to the expectations derived from the Dual Process Model. Moreover, gaze patterns over the course of the decision suggested that preference-consistent choices were made with less decision conflict.

In sum, this dissertation therein addresses a contemporary theoretical debate about the cognitive processes in moral decision making. In combination with the strongly related investigation of cognitive processes in the domain of prosociality in group settings, a research trajectory with a focus on the cognitive underpinnings of decision behavior is pursued.

Studying Cognitive Processes

Given the demonstrable interest of the scientific community and the use of process tracing techniques in this dissertation, here, a critical overview of methods that have been used to infer cognitive processes is presented. Examples from the domains of prosociality and morality are included for illustration.

Hypothetical and Incentivized Decision Settings

To study cognitive processes underlying respondents' decisions, participants must first be exposed to a decision setting. Two general tendencies for the selection of such settings are apparent in the literature: One stream of the literature uses *vignettes*, hypothetical situations which are designed to reveal participants' characteristics and judgments. Another stream of the literature uses *incentivized* decision situations, where participants' choices have real-life consequences.

In the domain of prosociality and morality, *vignettes* are frequently used. As a classical example for the use of vignettes to study prosociality, consider a study on the influence of fitness-related circumstances on helping behavior (Burnstein, Crandall, & Kitayama, 1994). In Study 3 of this paper, participants were asked to imagine they were a citizen of a sub-Saharan country suffering from famine and disease. Further, participants were asked to imagine that two members of their family (e.g., 75-year-old grandfather and 3-month-old nephew) were asleep in a burning building, and they only had resources to recue one of them. Given these circumstances, participants were asked to indicate whom they would prefer to help (for a more recent use of these vignettes, see Graziano, Habashi, Sheese, & Tobin, 2007). In the domain of morality, there is also a large body of research which relies on the use of moral vignettes (Christensen & Gomila, 2012; Cushman & Young, 2009) such as the trolley dilemma outlined above, immersing the respondent in hypothetical situations pitting different moral values against each other.

However, the use of vignettes to elicit choices relies on a number of auxiliary assumptions, which must be borne in mind when interpreting responses and inferring the validity of processes assessment. In particular, it is questionable whether participants can indeed imagine the situational context: It may be easier to imagine certain vignette scenarios than others, which could bias the process results obtained. Beyond participants' conception of the vignette context, the fact that the judgments are entirely hypothetical remains a problem. On the one hand, what participants think they would do might be a solid approximation for what they will do, assuming that they are able to accurately predict their own behavior and to report it accurately as well. While some believe that participants' intrinsic motivation in laboratory studies is high, leading them to invest enough effort to consider the outcomes of their decisions even in hypothetical situations (Camerer & Hogarth, 1999), others argue that non-hypothetical situations increase the effort participants invest in making their decisions, creating more naturalistic decision settings in the lab (Rosenboim & Shavit, 2012). To

determine if differences in respondents' behavior must be expected between hypothetical situations and situations where their decisions has real consequences, a number of studies in varying domains of behavioral economics can be consulted (for overviews, see Hertwig & Ortmann, 2001; Kühberger, Schulte-Mecklenbeck, & Perner, 2002). In the domain of social preferences, findings generally point to the interpretations that hypothetical vs. incentivized settings yield different choices.

Following the argument that incentivizing choices creates more externally valid decision settings in the laboratory (Rosenboim & Shavit, 2012), and assuming less noise in these measurements (Hertwig & Ortmann, 2001), using *decision situations with real consequences* has become increasingly popular in empirical social science generally, and in the study of prosociality and moral behavior specifically. The use of economic games to study prosociality is widespread, employing different decision problems such as the dictator games and prisoner's dilemma games introduced above. For example, a meta-analysis of behavior in dictator games addresses over 130 papers using this task to study generosity (Engel, 2011). SVO is also frequently studied via an incentivized measure of prosociality (for an overview, see Murphy & Ackermann, 2013). In the domain of moral decisions, the advent of incentivized settings has occurred only recently. In a seminal paper, participants were asked to decide between saving the life of surplus lab mice or receiving a payoff of 10€ in the individual decision treatment (Falk & Szech, 2013). When market forces were active through the introduction of a bi- or multilateral market, participants could forgo saving the life of a mouse for a price a buyer endowed with 20€ would offer. In these market scenarios, more participants were willing to sacrifice the mouse's life for less than or equal to 10€, indicating that market forces eroded moral decision making. In a follow-up study, donations to fund leprosy operations were used as a moral good while investigating the use of the replacement excuse (i.e., thinking that if one does not do X oneself, another person in the market will do X) depending on the presence of social norms (Bartling & Özdemir, 2017). In another example, a modified dictator

game was used to study people's willingness to harm an orphan by taking away a donation to purchase food in order to benefit others with a greater donation (Perera, Canic, & Ludvig, 2015).

In sum, both when investigating prosocial and moral decision making, using choices that have real consequences is the most promising setting to gain insights into people's preferences, decisions and decision processes.

Process Measures

To investigate cognitive processes, a number of potential methods can be considered. Participants could be asked how difficult it was to make their choice, or in which order they weighed the evidence for and against each option. Although this approach may be intuitively appealing, two important problems of such a self-report approach exist. First, can participants accurately assess how difficult it was for them to make the choice? Participants may lack insights into their own reasoning process (see the debate sparked by Nisbett & Wilson, 1977). Second, will participants accurately report their subjective decision difficulty? In particular in the domains of prosociality and morality, social desirability may bias participants'' responses. To avoid these issues, a variety of measures of cognitive processes that do not rely on introspection have been proposed. Here, an overview of a number of options for assessing cognitive processing during the decision without relying on self-reports is presented.

A frequently used and easy to employ option would be to measure decision times. Sometimes decision times are used as an indicator of intuitive vs. deliberate processing mode, assuming that intuitive decisions take less time than deliberate and effortfully reasoned choices (e.g., in the domains of morality and prosociality: Greene, Nystrom, Engell, Darley, & Cohen, 2004; Greene, Sommerville, Nystrom, Darley, & Cohen, 2001; Piovesan & Wengström, 2009; Rand et al., 2014, 2012).

Problematically, relying on reaction time data to infer whether a cognitive process is intuitive or deliberate can be subject to the reverse inference problem (e.g., in the domain of prosociality, see Krajbich et al., 2015). Although intuitive decisions have been shown to be faster (J. S. B. T. Evans & Stanovich, 2013), it is not necessarily true that fast decisions are intuitive. Other processes affecting response times could occur simultaneously, rendering the assumption that decisions are fast if and only if they are intuitive a potential logical error. For instance, Evidence Accumulation Models have argued that decision difficulty also affects response times in prosocial decision making (Krajbich et al., 2015). Therefore, in the absence of manipulations of intuitive vs. deliberate processing, reaction time data must be interpreted cautiously in the light of the auxiliary assumptions for its meaning.

A break-through in the study of decision processes was achieved with the notion that the reasoning process itself could be followed in real time. Rather than assessing how long a decision took, so-called process tracing methods were conceived to illuminate which reasoning steps decision makers made during the decision process. Here, a brief overview of different process tracing methods and their application is sketched.

Think-Aloud Protocols. One of the first process tracing methods were think-aloud protocols (Ericsson & Simon, 1993; Svenson, 1974). This method asks participants to verbalize their thoughts while completing a decision task. For example, in a study on moral reasoning, participants were confronted with vignettes in which different actors were depicted as behaving in more or less morally praise- or blameworthy ways (Bucciarelli, Khemlani, & Johnson-Laird, 2008, Study 3). While deciding who deserved more praise or blame, participants were asked to verbalize their reasoning process. Results suggested that decision makers were more likely to reason before making a choice than to choose immediately without reasoning. Problematically, deriving insights about cognitive processes from think-aloud

protocols requires the assumption that respondents can fully report their reasoning process. However, some evidence suggests that think-aloud protocols are detailed enough for computer programs fed with reported reasons to reach the same conclusions (Van der Henst, Yang, & Johnson-Laird, 2002). Nevertheless, this method remains problematic because responses must be coded and categorized, decreasing the advantages of this method.

Information Boards and Mouselab. Pursuing the idea to record appraisal of information, information boards were developed, where information about choice options remained hidden unless participants explicitly accessed it (Payne, 1976). This technique was later computerized such that participants could unveil information via mouse clicks or hovering over the area where information is hidden, and information would be covered again when they left the area (Mouselab; Bettman, Johnson, & Payne, 1990; Payne, Bettman, & Johnson, 1993), simplifying and increasing the objectivity of its use. In both cases, the steps of information search and the subsequent choice are recorded. Problematically, costs for accessing information are imposed, intervening in the information search strategy and subsequently changing choices compared to unperturbed reasoning processes (Glöckner & Betsch, 2008; Lohse & Johnson, 1996). Although an important breakthrough for process tracing research in decision making (Beach & Potter, 1992), these methods therefore carry marked limitations.

Mouse Tracking. Building on the premise of inferring decision processes from computer mouse movements, mouse tracking was developed to trace the curvature of mouse movements towards specific targets (Dale, Kehoe, & Spivey, 2007; Spivey, Grosjean, & Knoblich, 2005). For example, in an application to the domain of prosociality (Kieslich & Hilbig, 2014), participants were confronted with different social dilemmas, and choices were made by clicking one of two buttons symbolizing cooperation or defection, and the mouse trajectories towards these targets were

analyzed. When choosing to cooperate, the trajectory was less curved towards the non-chosen option, indicating less cognitive conflict in cooperation.

In sum, although important insights were accomplished using the process tracing techniques introduced above, they have critical limitations. The methods both interrupt and fail to capture the decision process on a sufficiently fine-grained level of analysis. Therefore, many questions regarding the cognitive processes driving prosociality and morality were left unanswered.

Eye Gaze and Decision Processes

To address the need for a fine-grained and unobtrusive measure of cognitive processes in the decision process, eye tracking has emerged as an increasingly popular tool. This section gives an overview of the rationale of eye tracking as a tool to study cognitive processes in general and in decision processes in particular. A short introduction to the technical functioning and possible dependent variables is given and related to predictions of general process theories of decision making. Finally, a short overview of the use of eye tracking in the domains of prosociality and morality is provided.

The Link between the Eye and the Mind. Eye gaze is used to study selectivity in perception, which is defined as visual attention (for an overview of the perception-attention link, see Orquin & Mueller Loose, 2013), to infer how information acquisition and cognitive processing take place (Ashby, Johnson, Krajbich, & Wedel, 2016). This strategy relies on the Eye Mind Hypothesis (Just & Carpenter, 1980), positing that what is visually fixated on is being processed in working memory ("There is no appreciable lag between what is fixated and what is processed."; Just & Carpenter, 1980, p. 331). The human eye can monitor a visual field of about 200° with varying degrees of clarity, but the fovea, the point of clearest vision, only covers a 2° visual angle (Levi, Klein, & Aitsebaomo, 1985). Rapid eye movements called saccades are

used to position the eye so that the reflection of stimuli enters the foveal region. During fixations, instances where a stimulus is held in the foveal region for a relatively long period of time (more than 50ms; Salvucci & Goldberg, 2000; but typically around 200ms to 300ms; Holmqvist et al., 2011, p. 21), focused stimuli can be processed. Sometimes information that is not held in the foveal region, but in neighboring regions of the eye where vision is less clear, can still be processed (for examples of experts' processing of parafoveally regarded stimuli, see Charness, Reingold, Pomplun, & Stampe, 2001; Reingold, Charness, Pomplun, & Stampe, 2001). Critically, information cannot be processed when it is passed over in a saccade (Rayner, 1998), because the speed at which the eye moves during saccades is too high for information to be detected and processed (saccadic suppression; Matin, 1974).

At the same time, there are also indications that visual attention does not necessarily indicate that certain stimuli are being cognitively processed, contradicting the eye-mind hypothesis. For instance, visual attention and awareness are not identical (Lamme, 2003), and attention can shift even when the processing of a certain piece of information has not been completed (Engbert, Nuthmann, Richter, & Kliegl, 2005). Moreover, eye gaze and attention can also be decoupled (Posner, 1980): When attention is covert (as opposed to overt), people may look at one stimulus but think about something else. These caveats have led to the formulation of less strong versions of the original Eye Mind Hypothesis in which a one-to-one relationship of attention and cognitive processing is no longer assumed. One such reformulation is that of Huettig, Olivers and Hartsuiker (2011, p. 141): "The most active location in working memory will eventually determine the most likely direction of the eye movement at a given point in time.".

This reformulation further points to the directionality of the relation of eye gaze and attention: Does eye gaze determine what will be processed, or will eye gaze follow cognitive shifts in attention? Both directions are argued to occur. Shifts in attention

brought about by an involuntary reorientation are referred to as exogenous, bottom-up or stimulus-driven attention shifts, while endogenously determined, top-down or goal-directed attention shifts describe the voluntary orientation towards a certain stimulus (for overviews, see Corbetta & Shulman, 2002; Theeuwes, 2010; Yantis, 2000). A number of stimulus features can determine the *exogenous direction of attention* (for a review, see Orquin & Mueller Loose, 2013). An early theory on stimulus-driven attention suggested that the visual field could be perceived as a saliency map, where features such as luminance or direction of movement could increase the likelihood that an area captures attention (Koch & Ullman, 1985) For other examples of stimulus features capturing attention, Lohse (1997) showed that colored or large stimuli received more attention. Moreover, the position of stimuli relative to others has been shown to influence attention: stimuli displayed in the center (central position effect; Chandon, Hutchinson, Bradlow, & Young, 2009) or at the top (Sütterlin, Brunner, & Opwis, 2008) receive more attention. Overall, research shows that such saliency-based models perform better than chance in predicting attention shifts (Foulsham & Underwood, 2008).

Top-down shifts in attention are famously illustrated with the rubber band metaphor (Holmqvist et al., 2011, p. 379): When part of a rubber band is stretched to one point (i.e., when attention is directed at a certain piece of information), the other parts of the band will follow (i.e., eye gaze will be directed to the relevant stimulus). When attention is endogenously redirected, the lag between this shift in attention and the subsequent shift in eye gaze is about 250ms (Deubel, 2008). When making decisions, gazes can be categorized into three types of distinct functionality (Glaholt & Reingold, 2011): initial short fixations to gain an overview (Glöckner & Herbold, 2011) and decide which alternatives to incorporate into the set of options considered in depth (Russo & Leclerc, 1994; Wedell & Senter, 1997), longer fixations in the middle of the decision process to compare the alternatives (Russo & Leclerc, 1994), and

finally shorter fixations to confirm the emerging decision tendency (Krajbich, Armel, & Rangel, 2010; Krajbich, Lu, Camerer, & Rangel, 2012; Krajbich & Rangel, 2011).

In a seminal investigation of how different processing goals determine eye gaze, scan paths were compared when participants repeatedly viewed artworks when given different instructions (Yarbus, 1967). How the painting was viewed depended dramatically on which goals participants pursued, such as when estimating the material wealth or assessing the age of a depicted family, and when trying to memorize the location of people and objects in the painting (e.g., Yarbus, 1967, p. 174). Since then, a large body of research has been devoted to studying how top-down attention shifts occur (for reviews, see Baluch & Itti, 2011; Noudoost, Chang, Steinmetz, & Moore, 2010). Here, a special focus is placed on investigations of top-down influences on attention in decision making research (for an overview, see Orquin & Mueller Loose, 2013). A number of studies have investigated how different decision tasks affected information search patterns. For instance, time pressure to make a decision quickly resulted in altered scan paths and decreased the number of features fixated on (Pietersa & Warlopb, 1999). In risky choices, preference reversals between different choice tasks (Lichtenstein & Slovic, 1971) are accompanied by diverging gaze patterns (B. Kim, Seligman, & Kable, 2012): When people decided between two risky gambles, they focused more on probabilities and subsequently were more likely to choose the gamble with the more probable outcome but lower expected payoff. However, when betting on a single risky gamble, outcome amounts were attended more and higher outcome gambles with low probabilities received higher prices than high probability gambles.

The latter example also alludes to the relation of *utility* and eye gaze: Decision makers are more likely to gaze at information that has a high utility (Fiedler & Glöckner, 2012) or is important to the decision maker (Glöckner, Fiedler, Hochman, Ayal, & Hilbig, 2012; Glöckner & Herbold, 2011; van Raaij, 1977). People fixate more on the

decision alternative they will later choose (Russo & Leclerc, 1994), and are more likely to direct their last fixation before deciding to the later chosen alternative (Krajbich et al., 2010, 2012; Krajbich & Rangel, 2011). This attention shift is not a sudden change, but is reflected in a gradual increase in fixations on the later chosen option over the course of the decision process (Gaze Cascade Effect; Shimojo, Simion, Shimojo, & Scheier, 2003; Simion & Shimojo, 2007). Moreover, decisions that are more difficult, i.e. where the alternatives are very similar in their respective utility, are associated with higher fixation counts (Fiedler & Glöckner, 2012; Glöckner & Herbold, 2011; Krajbich et al., 2010, 2012; Krajbich & Rangel, 2011).

Currently, a debate is ongoing about the question whether *attention can bias choices*. This idea that attributes which receive more attention bias the decision was formulated both in the literature on gaze cascade effects (Shimojo et al., 2003) and in the attentional drift diffusion model (Krajbich et al., 2010). There have been successful attempts to bias simple choices by manipulating visual attention (Armel, Beaumel, & Rangel, 2008; Milosavljevic, Navalpakkam, Koch, & Rangel, 2011) but the strong formulation of the causal role of attention on decisions in more complex domains such as morality (Pärnamets et al., 2015; but see Newell & Le Pelley, 2018) remains to be shown. Rather, empirical evidence is argued to be more consistent with a form of mere exposure effect (Glaholt & Reingold, 2011; Glaholt, Wu, & Reingold, 2009; Shimojo et al., 2003), i.e., an increased likelihood to choose an alternative to which one has been exposed longer, in combination with an increased likelihood to choose the more salient alternative based on bottom-up effects (e.g., Chandon et al., 2009; Lohse, 1997).

In sum, the relation of eye gaze and attention, and ability to illuminate which information is being processed and how by investigating gaze behavior have been well established in the literature.

Measuring Eye Gaze. Given that eye gaze can be considered as a helpful tool for the investigation of cognitive processes underlying decision making, the following section gives an overview of the technical side of measuring eye gaze.

Eye tracking methodology has its roots in reading research at a time where eye motions were first investigated using mirrors, leading to the initial observation of saccades (Javal, 1878). Moreover, saccades were counted while participants read with one eye by counting the bumps of the other, closed eye's cornea against a microphone (see Tinker, 1928). Since then, eye tracking tools have undergone an extreme development (for detailed overviews, see Richardson & Spivey, 2008; Wade & Tatler, 2005), passing a stage where a writing device was connected to the participants' eyes via wires and contact lenses (Delabarre, 1898; Huey, 1898), and ultimately evolving to contemporary techniques using video recordings of the eye.

To identify the location to which gaze is directed (point of regard) independently of movements of the head, these recording tools base real-time calculations on corneal reflections relative to the center of the pupil (for an in-depth description, see Duchowski, 2007, Chapter 5.4). An infrared light ray is directed at the eye and reflected by the cornea in four light rays (Purkinje reflections, see Crane, 1994). In a calibration phase, the position of the first Purkinje reflection relative to the center of the pupil is assessed while the participant is directed to gaze at specific locations. Subsequently, when participants freely gaze at different locations, their point of regard can be calculated.

Dependent Variables. A number of variables is used to characterize eye gaze based on fixations, saccades, blinks and the pupil. Here, key variables are introduced and related to the cognitive processes that have been inferred from them in the existing literature.

First, measures related to *fixations* have already been referred to above. Three attributes of fixations are used to infer cognitive processes: their location, their frequency and their timing. The location on which decision makers fixate, or rather the information symbolized in these locations, is used to infer attention, and the content of information processed in working memory. Often, certain stimuli are embedded in so-called Areas of Interest (AOIs), which are slightly larger than the stimuli themselves to account for inaccuracies in the gaze recording (for an in-depth discussion, see Orquin, Ashby, & Clarke, 2016). From counting the *number of fixations* (also called *fixation count*) inside an AOI, it is inferred how much attention decision makers spend on the stimulus. Further, the number of fixations is used to quantify the relative importance or weight of the contained piece of information: more important pieces of information are fixated more often (Russo & Leclerc, 1994). In Chapters 2 and 3, we use the number of fixations as an indicator for decision effort, relying on the auxiliary assumption that higher effort is required to resolve decision problems of high difficulty, which in turn increase the number of fixations (Krajbich et al., 2010). A second indicator for decision effort is the number of pieces of information fixated on out of all available pieces of information, relying on the assumption that a more effortful decision process integrates more pieces of information (Fiedler, Glöckner, Nicklisch, & Dickert, 2013). The number of pieces of information fixated on as an indicator for decision effort is used in Chapters 2 and 3. Another variable of interest related to fixations is the order in which fixations to specific AOIs occur. One way of utilizing fixation sequences is studying gazes to the choice options over time. Differences in how fixations are allocated to two available alternatives as the decision process evolves can be assessed to infer cognitive conflict: the larger the difference in fixations over time, the smaller the cognitive conflict occurring between the competing options. This measure is utilized in Chapter 4. Another way of using fixation sequences is by constructing full profiles of gazes to all AOIs called *scan paths*. Using scan paths makes it possible to infer patterns of attention and search strategies to understand

the sequence in which information is attended to and processed (Noton & Stark, 1971). For instance, in multi-attribute choice where choice options are described by several characteristics, attribute-wise vs. alternative-wise search patterns calculated as a so-called search index could be used to infer heuristic search strategies (Payne, 1976). In another example, the Needleman-Wunsch algorithm (Needleman & Wunsch, 1970) was applied to compare scan paths in different risky choice tasks (Day, 2010). Regarding the timing of fixations, the *total dwell time*, i.e., the total duration of all fixations directed to the AOI, can be used to infer processing elaborateness (Velichkovsky, Rothert, Kopf, Dornhöfer, & Joos, 2002). Another variable of interest is the duration of single fixations, that is how long eye gaze "rests" on a specific location. Fixation durations are used to infer depth of processing (Velichkovsky, Dornhoefer, Pannasch, & Unema, 2000; Velichkovsky et al., 2002). Short fixations indicate information scanning and lower-level automatic processing (Glöckner & Herbold, 2011), while higher-level deliberate processing prolongs fixation durations.

Second, other dependent variables make use of *saccades*. For instance, the *saccadic distance*, i.e., the distance between two fixations, is used to infer local vs. global attention states (Lans, Pieters, & Wedel, 2008; Wedel, Pieters, & Liechty, 2008). Global attention is used to orient oneself in the visual environment and to find a target stimulus, for which longer saccadic distances are used. When switching to local attention and shorter saccadic distances, a target stimulus is processed in detail. *Saccadic velocity*, the speed of the eye motion between two fixations, is argued to indicate cognitive load: Higher cognitive load was associated with faster saccades (Bodala, Ke, Mir, Thakor, & Al-Nashash, 2014). Relatedly, evidence indicates that decision-irrelevant saccades were faster when decision difficulty was higher due to incongruent or inconsistent stimuli (Joo, Katz, & Huk, 2016).

Third, *eye blinks*, that is the rapid opening and closing of the eyelid, can be counted over time to constitute the *blink rate* (number of blinks per minute). Although blinks are a biological necessity to maintain the moistness of the eye, the blink frequency has also been argued to depend on processing difficulty (Stern, Boyer, & Schroeder, 1994).

Fourth, dependent variables make use of the size of the *pupil*. The pupil can be dilated or constricted, which is a biological process regulating the amount of light admitted to the cornea. Beyond the biological response to different lighting and luminance conditions, *pupil dilation* also depends on cognitive processes. It is chiefly used to infer arousal and cognitive load (Andreassi, 2010; Beatty & Lucero-Wagoner, 2000; Kahneman & Beatty, 1966).

In sum, eye tracking offers a versatile toolbox of variables that can be used to operationalize specific cognitive functions. Additionally, the use of multiple related dependent variables is especially advantageous to perform precise tests of theoretical process predictions (Glöckner & Betsch, 2011).

Theories of Decision Processes

The theories of prosocial and moral decision making introduced in the beginning make some explicit process predictions, but few predictions are directly specified for dependent measures that can be gained from eye tracking. Therefore, more general theories of decision processes could be consulted here, which describe the role cognitive processes are assumed to play during the decision making process (also see Marewski, Bröder, & Glöckner, 2018) and make specific predictions for eye tracking variables. One of these general theories is reviewed here to provide context for the specific process models referenced above, and as the basis for process predictions specific to eye tracking which are derived in the subsequent chapters of this dissertation.

Evidence Accumulation Models of decision making, such as the Decision Field Theory (Busemeyer & Townsend, 1993), as well as the *Attentional Drift Diffusion Model* (aDDM; Krajbich et al., 2010, 2012; Krajbich & Rangel, 2011) share the notion that that information is accumulated in a sequential sampling process until a threshold is passed and a choice is made. Utility or decision value of the available options determine the choice. One of the key differences between the decision field theory and the aDDM is that the former assumes that attention is distributed following a weighted additive function, which is consistent with most empirical evidence on changing fixation likelihoods depending on task characteristics and utility, while the latter assumes a stochastic distribution of attention. The aDDM yields a large number of specific process predictions. It posits that attention to an option increases its relative value, indicating that the option gazed at first is initially favored over the other alternatives, and that any mechanism directing more attention to one alternative will increase the likelihood that this option is chosen. Moreover, the rate at which the decision threshold is approached is argued to be steeper the larger the utility differences between the alternatives. The model further proposes a stable fixation duration across the decision process that depends on the choice difficulty (in contrast to empirical evidence suggesting changed fixation durations depending on characteristics such as time pressure (Pietersa & Warlopb, 1999) and utility (Fiedler & Glöckner, 2012), with the exception of the last fixation, which is argued to be cut short as the decision threshold is passed. This also means that the last fixation is assumed to be allocated to the later chosen alternative.

Given its relative parsimony in assumptions about the decision process, and the formulation of highly specific processing predictions, the evidence accumulation models and the aDDM in particular have emerged as a frequently employed decision theory. In the scope of this dissertation, we make use of the predictions of the aDDM to derive specific hypotheses in the context of prosocial choices (Chapter 3), and to

motivate measures used to investigate moral decision making, we rely on auxiliary assumptions derived from findings in the context of the aDDM (Chapter 4).

Eye Gaze in Moral and Social Decision Making. Finally, some prior research has made use of eye gaze to study decision making in the domains of prosociality and morality, which is briefly summarized here to provide context for the subsequent empirical contributions made in the scope of this dissertation.

Regarding prosociality, a recent paper on the processes underlying prosociality showed that decision makers' SVO determined how information was searched for in Dictator Games and Public Goods Games (Fiedler et al., 2013). With increasing deviation from prosocial SVO, increases in decision times, fixation counts and the proportion of inspected information, as well as the proportion of attention to others' payoffs and transitions between own and others' payoffs were found. These findings point to a weighted additive decision strategy in prosocial decision making, in line with Evidence Accumulation Models. The findings also directly contradicted the prediction that cooperation is intuitive derived from the Social Heuristics Hypothesis. Another study supported the notion that social preferences are distinctly related to the proportion of attention allocated at information relevant for the subsequent choice (Jiang, Potters, & Funaki, 2015). Moreover, another study has contributed eye tracking evidence in group settings, albeit not directly addressing prosociality but trust. Pupil dilation was used to study behavior in incentivized trust games (Kret, Fischer, & De Dreu, 2015). Trustees were trusted more when they had dilated pupils, and trustors' pupils mimicked trustees' pupil dilation. Trust was predicted by trustors' mimicking of ingroup members' pupil dilation. Finally, other work has demonstrated that ingroup members receive preferential visual attention (Kawakami et al., 2014).

Regarding the relation of attention and morality, in a study using trolley-type dilemmas post-decisional attention was shown to be directed away from the victim, suggesting a motivation to avoid feelings of guilt (Kastner, 2010). Moreover, a number

of studies have addressed the processes underlying moral decision making via eye tracking in contexts other than trolley-type dilemmas (for an overview, see Fiedler & Glöckner, 2015). For instance, recent work using pupil dilation has found that the dilation of the pupil predicted deceptive messages to a partner (Wang, Spezio, & Camerer, 2010) and dishonest reporting of outcomes in a simple counting task (Hochman, Glöckner, Fiedler, & Ayal, 2016) to increase one's own payoffs.

Overall, eye tracking has been shown to be a useful tool to investigate, in real-time and unobtrusively, the cognitive processes underlying decision making. The use of various dependent variables enables a variety of insights into information search behavior, cognitive effort, and the importance of certain pieces of information for the decision maker. Finally, initial evidence suggests that eye tracking is well suited for the investigation of decision processes in the domains of prosociality and morality.

Summary

In sum, this introduction has given an overview of the historical and current debates in research on prosociality and morality, finding parallel trends in their development, such as with regard to the increasing use of decisions with real consequences as a paradigm for investigation, or the concurrent juxtaposition of dual process models and choice differentiability accounts. Importantly, in both areas of research, a growing interest in cognitive processes is observed, which this thesis addresses. In particular, the cognitive processes underlying prosocial decision making in group settings and moral decision making are identified as areas where further research is needed. This need for further empirical investigations in the domain of prosociality specifically concerns the acquisition of information about others' group membership, attention towards own and others' outcomes, and effort invested in information search during the decision process. In the domain of moral decision making, we test the predictions made by competing theoretical models, the dual-process theory of moral decision making, and the choice discriminability account,

about the relation of deontological vs. utilitarian decision preferences and decision effort. In a discussion of empirical research methods potentially usable for investigating cognitive processes, eye tracking is demonstrated to be the optimal methodology to conduct unobtrusive and fine-grained analyses of decision processes, and the aDDM is introduced as a general theory of decision processes, from which we derive predictions for the subsequent empirical investigations.