Learning together: behavioral, computational, and neural mechanisms underlying social learning in adolescence

Westhoff, B.

1

Chapter 1

General Introduction

The scope of this thesis

When you think about things you have learned as a child or teenager, you will likely think about particular skills or something you learned from (school) books, such as reading, math, or riding a bike. However, humans are sophisticated social beings who live in a dynamic and complex social environment, and encounter numerous social situations daily. Therefore, a considerable part of what we learn is *social* information, enabling us to navigate the social environment. For example, social interactions require quick learning about the behavior of other people. For instance, when you think about the first day of high school, you may recall this as an exciting day in which you were going to meet many new people. You did not know (most of) them yet: who were you going to like or dislike? Who would be your friends? Who would be your ideal project buddy? You had to learn about these people, which involves predicting others' behavior and interacting successfully with them. These are important elements of *social learning*.

Social learning thus involves predictions about other people – which affect your upcoming choices or actions – and these predictions will be updated when they do not match the actual outcome. Social learning may take different forms, which differ in the information one learns and its purpose. Historically, social learning theory was proposed by Albert Bandura (Bandura, 1977), which defines social learning as learning without direct reinforcement, such as learning *from* others by observing, imitating, and modeling their behavior, or vicariously learning from others' actions and outcomes (rewards and punishments). In addition, as mentioned above, social learning can involve learning *about* other people, such as whether you can or cannot trust someone. Finally, social learning may also encompass learning *for* others. That is, many of our actions may affect other people and we have to learn which actions we should repeat to help others. For example, to comfort an upset friend, we need to know what actions cheer them up. Learning to benefit or help others can also be referred to as *prosocial learning* (Lockwood et al., 2016). The current thesis focusses on learning *about* and learning *for* others. Although these two forms of social learning differ, they have overlapping elements and generally encompass learning about actions and their outcomes which involve other people.

Being able to learn adaptively in a social context is an essential social skill (Fareri et al., 2020). Social learning skills enable you, for instance, to know who you can/should cooperate with (e.g., in school or business projects), who you can trust with your secrets, and how best to help someone. Thus, the ability to learn about and for others helps you make appropriate choices that involve others, and determine how to behave around other people. Adaptive social skills are vital for building and maintaining healthy social relationships (Fareri et al., 2020), which have been shown to be beneficial for one's long-term wellbeing (Güroğlu, 2021; Uchino, 2009).

Social learning is a complex social skill that develops and improves across development. A developmental period that has been suggested to be particularly key for developing social learning skills, is adolescence (Blakemore & Mills, 2014; Crone & Dahl, 2012). Adolescence is a sensitive period for social development in general, and is characterized by large changes in the social environment. That is, adolescents spend more time with peers than with their family, and their social relations become more intense and complex (Brown & Larson, 2009; De Goede et al., 2009; Lam et al., 2014). Moreover, their sensitivity to social stimuli is heightened, which may, for example, result in increased susceptibility to peer influence (Blakemore & Mills, 2014; Crone & Dahl, 2012; Foulkes & Blakemore, 2016). These psychosocial changes are thought to result from ongoing structural and functional changes in the brain (Blakemore & Mills, 2014; Nelson et al., 2016; Somerville, 2013). Due to these environmental and neurobiological changes, adolescence may be a life phase particularly attuned to social learning.

In this thesis, I aim to examine the development of two forms of social learning across adolescence. Here, I will focus on (1) learning *about* other people, specifically, whether they are (un)cooperative and (un)trustworthy, and (2) learning *for* other people (prosocial learning) to know what actions may help others. To this end, I make use of multiple experimental paradigms in samples spanning early adolescence to late adolescence. A second aim was to examine underlying mechanisms and factors that account for age-related and individual differences in social learning. To this end, I combined self-report questionnaires, computational modeling, and functional neuroimaging.

In this first chapter, I first provide an introduction about adolescence as a key period for social development. Here, I discuss two social behaviors - specifically, trust behavior and prosocial behavior - that play an important role in social interactions as well as in the different forms of social learning. Next, I present reinforcement learning as a framework for social learning. Finally, I highlight the added value of using economic games, computational modeling, and functional neuroimaging for studying social decision-making and learning, and end with an outline of the empirical chapters.

Adolescence as a key phase for developing social learning skills

Adolescence is the developmental phase that marks the transition from childhood to adulthood. Adolescence has a biological starting point with the onset of puberty and accompanying hormonal and biological changes (+/- 9-12 years old) (Spear, 2011). The end of adolescence is culturally determined, but in Western cultures this is generally when one reaches mature social goals and is relatively independent of their parents. Thus, adolescence roughly spans 9-24 years of age, although different ages or (sub)labels are being used (Sawyer et al., 2018).

Adolescence is characterized by biological changes in levels of pubertal hormones, physical characteristics, and brain anatomy and function (see e.g., Dumontheil, 2016). Besides these biological changes, however, this life phase is also a time of profound psychological and social changes. Adolescence is a time of social reorientation, during which the social world and peer interactions become increasingly important. For instance, adolescents spend more time with peers than in childhood (De Goede et al., 2009), both offline and online (Lam et al., 2014), and they form and maintain high-quality friendships (Brown & Larson, 2009). Concurrently, there is a heightened sensitivity to social stimuli (Foulkes & Blakemore, 2016; Somerville, 2013) which may, for example, result in increased attention to peer evaluation (Guyer et al., 2014) and an increased susceptibility to peer influence (Blakemore & Mills, 2014; Crone & Dahl, 2012; Foulkes & Blakemore, 2016). Moreover, both the quantity and quality of social interactions change, and for functioning as an adult and achieving mature levels of social competence, it is essential that adolescents develop adaptive social skills. These are needed for e.g., building reciprocal social relations, which have been shown to be essential for long-term health and (emotional) wellbeing (Baumeister & Leary, 1995; House et al., 1988). In sum, adolescence is a key life phase for developing well-adjusted social behavior.

Well-adjusted social behavior comprises of multiple adaptive social learning skills. Next, I will discuss trust behavior and prosocial behavior, as these social behaviors play an important role in social interactions, as well as in learning about and learning for others, respectively.

Trust

A crucial component for cooperative social interactions and mutually beneficial interpersonal relationships, is trust. Trust is important at all levels of society – from interpersonal to institutional trust – and well-adjusted social behavior therefore also entails being adaptive with regard to deciding whether or not to trust. That is, trusting others who will reciprocate your trust will result in positive social interactions. However, trusting others who will betray your trust, will be wasteful for your resources. Showing maladaptive trust decisions, i.e., both trusting too often or too little, is likely to result in problems with peer relations and social behavior (Rotenberg et al., 2005). Therefore, it is vital for successfully navigating the social environment to be able to learn whom you can or cannot trust.

For studying social decision-making behaviors such as trust, economic games are well-suited. Economic games allow studying complex processes in a controlled experimental setup, and are moreover suitable for studying developmental patterns across large age ranges (Camerer, 2003; Gummerum et al., 2008). A well-known economic game for studying trust, is the trust game. In the original trust game, participants may have the role as either an investor or a trustee (Berg et al., 1995). As the investor, a participant may have two options: trusting (i.e., investing) or not trusting (i.e., not investing) the trustee. When choosing to trust, the investment (e.g., coins or tokens) will be multiplied with a certain factor, and the trustee can decide how much to reciprocate to the investor. Choosing to trust the trustee may result in the highest possible outcome in this game, but only if the trustee will actually behave trustworthy (i.e., reciprocate). Choosing not to trust the other will result in a guaranteed outcome, yet lower than a reciprocated trust choice. Choosing not to trust is optimal when the trustee is indeed untrustworthy, but choosing not to trust someone that turned out to be trustworthy, will result in a suboptimal outcome. Previous research has used (variations on) the trust game for studying developmental patterns of trust behavior across adolescence. These studies reported increasing levels of trust behavior from early- to mid-adolescence (e.g., (Fett, Gromann, et al., 2014; Sutter & Kocher, 2007; van den Bos et al., 2010; van den Bos, van Dijk, et al., 2012).

In the current thesis, chapters 2-4 use experimental paradigms that are based on these trust game principles to study adolescents' ability to learn about the trustworthiness of others (chapters 2 and 3), and, more specifically, how they sample information on others' trustworthiness in order to decide whether to trust someone (chapter 4).

Prosocial behaviors

Another form of social behavior that plays an important role in social relations, is prosocial behavior. Prosocial behaviors are defined as voluntary social behaviors that are intended to benefit others (Padilla-Walker & Carlo, 2014). It is a multidimensional construct that involves a wide variety of behaviors, such as helping others, sharing resources, cooperating, and comforting. Studies have shown that prosocial behavior is important for being liked by others, and building and maintaining positive reciprocal social relationships (Güroğlu et al., 2007; Peters et al., 2010). With regard to developmental patterns of prosocial behavior, studies have shown that across age, adolescents increasingly exhibit prosocial behaviors in the domain of sharing and giving (Güroğlu et al., 2014; Padilla-Walker & Carlo, 2014).

To be able to show prosocial behaviors, it is essential that you are able to learn which of your actions have beneficial consequences for the other person (i.e., prosocial learning). Previous studies investigating prosocial learning, have used prosocial learning tasks in which participant could learn to obtain rewards for others (prosocial learning), and distinguished that from self-benefitting learning (Cutler et al., 2021; Lockwood et al., 2016; Sul et al., 2015). Although numerous developmental studies on prosocial behavior have been conducted in the past years (e.g., Güroğlu et al., 2014; Schreuders et al., 2018; van de Groep et al., 2020; van Hoorn et al., 2016), no studies have investigated prosocial *learning* across development. In chapter 5, I used such a prosocial learning task to investigate age-related differences in adolescents' abilities to learn for others compared to learning for self.

Reinforcement learning framework applied to social learning

A framework of interest to understand the mechanisms of social learning, is reinforcement learning (Figure 1A). Reinforcement learning enables describing how decisions and their outcomes are paired over time (Sutton & Barto, 1998). This reinforcement learning framework can be captured in a mathematical framework, which has been widely applied in many areas of psychology and neuroscience (Dayan & Balleine, 2002; Zhang et al., 2020). Reinforcement learning is a process in which our past experiences are used to perform actions that are likely to result in positive outcomes. Reinforcement learning depends largely on *prediction errors* (PEs), which reflect the difference between the expected outcome and actual outcome of an action. When the outcome is better or worse than expected, there is a PE, and its size depends on how large the deviation is. PEs drive learning as they are used to update the future expectation with the new information. That is, an outcome that is better than expected (positive PE), will increase our expectations that performing that action again will result in a positive outcome, and we will become more likely to repeat the action in the future. When an outcome is worse than expected (negative PE), we update our expectations as such that we will be more likely to avoid that action in the future. The extent to which expectations are updated is determined by a *learning rate*. For someone with a relatively high learning rate, the weight of the PE is larger and their expectations are being updated to a greater extent. For someone with a rather low learning rate, however, the weight of a PE is lower and their expectations are updated only to a limited extent. As a result, they incorporate the feedback not only from the last situation but from a longer time frame, which could e.g., be beneficial in relatively stable learning contexts (Nussenbaum & Hartley, 2019).

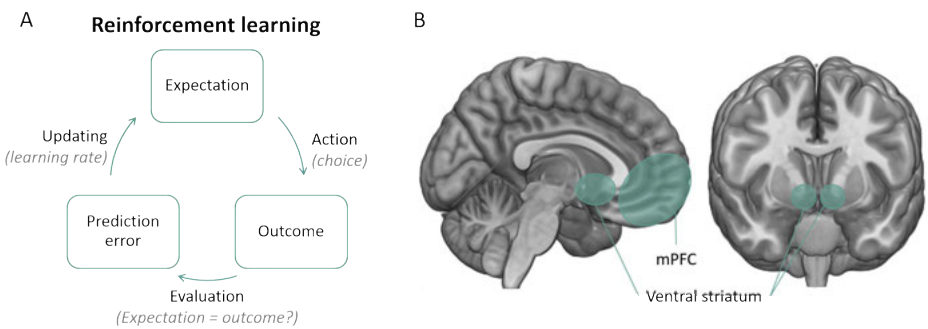

Figure 1. (A) Reinforcement learning framework. The action you perform is based on your expectation of the resulting outcome. During the evaluation, the actual outcome of your action is compared to your expected outcome. The difference between these, is the prediction error. A prediction error of 0 indicates that the outcome precisely matches your expectations. If the actual outcome differs from your expected outcome (prediction error ≠ 0), this will be used to update your expectations for the

subsequent action. The extent to which you update your expectations is determined by your learning rate. **(B)** Ventral striatum and medial prefrontal cortex (mPFC) are core brain regions involved in reinforcement learning.

For using the reinforcement learning framework to investigate (social) learning, researchers make use of computational modeling. Computational models are mathematical models applied to behavioral data. The advantage of such a computational model is that they can describe behavior, and as such can identify underlying mechanisms that you cannot directly observe or measure from behavior (i.e., latent variables) (van den Bos et al., 2018). Therefore, computational modeling is a valuable tool for studying people's behavior and the underlying mechanisms of individual or developmental differences therein.

A novel perspective is to apply this reinforcement learning framework to social learning, as it is thought to rely, at least partly, on similar computational learning mechanisms (Joiner et al., 2017; Lockwood & Klein-Flügge, 2020; Zhang et al., 2020) and potentially also similar neural mechanisms (Lockwood et al., 2020) as those used in basic reinforcement learning. For example, computational reinforcement learning models in social learning contexts have assessed whether learning rates differ between learning conditions (e.g., learning to benefit oneself compared to learning to benefit others (Cutler et al., 2021; Lockwood et al., 2016). However, in social psychology and neuroscience, variations on these reinforcement learning models may be insightful as well (Zhang et al., 2020). For instance, these reinforcement learning models can be extended with social elements, such as social preferences (inequality aversion), to better capture the underlying processes involved in social learning.

In this thesis, I applied computational modeling to study how adolescents learn *about* and *for* other people, and examined developmental and individual differences in the underlying computational mechanisms of social learning.

Social reinforcement learning in the developing brain

An additional valuable methodological approach for understanding underlying mechanisms of a particular behavior, is the use of functional neuroimaging. When studying the brain, researchers often make use of Magnetic Resonance Imaging (MRI) scanners. This noninvasive method is suitable to use in participants from a relatively young age. When the MRI scanner is used to measure the blood-oxygen-level dependent (BOLD) signal in the blood vessels in the brain, this is called functional MRI (fMRI). These oxygen levels are an indirect measure of neural activity: when neurons in a certain brain area are active, they will need oxygen, thus higher oxygenated levels indirectly indicate more activity in that particular brain area. When participants perform a task (e.g., a social learning task) during fMRI, the scanner takes images every few seconds. This technique allows researchers to study which brain areas are involved during which parts of the tasks. fMRI has a good spatial resolution, which allows locating the BOLD signals during specific tasks or events.

fMRI has been used across an extensive range of experimental paradigms. For studies on (social) learning fMRI can add insights in what neurobiological mechanisms underlie reinforcement learning. When this is combined with computational modeling, this allows for more sophisticated probing of the learning processes. That is, the PEs (representing learning signals) are extracted on a trial-by-trial basis from the computational models. These PEs can be used in the fMRI analyses such that the learning signals can be tracked in the brain. Key regions that have been linked to reinforcement learning are the ventral striatum and the medial prefrontal cortex (mPFC) (see Figure 1B). The striatum is a region located in the mid-brain that receives dopaminergic input from the ventral tegmental area and substantia nigra, and is functionally related to reward processing and (social) learning (Olsson et al., 2020). The striatum has connections with the mPFC. The mPFC is thought to integrate reward (value) and is involved in (social) decision-making and reinforcement learning (Joiner et al., 2017). Although social learning may depend on several interacting brain regions, these reward-related regions are expected to be key for social learning.

Thus, using fMRI for studying social learning provides insights into its underlying mechanisms. In chapter 5, I combined fMRI with computational modeling of prosocial learning data, and examined age-related differences in prosocial learning signals in the ventral striatum and mPFC.

Outline of this thesis

In this thesis, I report the results from four empirical studies that I have conducted to investigate the development and underlying mechanisms of social learning in typically developing adolescents, using experimental behavioral paradigms, self-report questionnaires, computational modeling, and fMRI.

In **chapter 2**, I present an experimental paradigm consisting of multiple 1-shot economic games to examine adolescents' ability to learn *about* and adjust to others that differed in their levels of cooperative behaviors (i.e., trustworthiness and cooperation). More specifically, in one condition ('trust game') participants had to learn to trust trustworthy others and not trust untrustworthy others. In the other condition ('coordination game'), participants had to learn to coordinate their choices by accepting either an advantage from cooperative others or a disadvantage from uncooperative others. In a large adolescent sample spanning a broad age range (N = 244, 8-23 years), I examined age-related differences and factors (e.g., learning rates) underlying individual differences in participants' ability to learn to adjust to others' behavior.

Chapter 3, describes an experimental study in an adolescent sample with a broad age range (*N* = 157, 10-24 years) that focused on the trust game condition from chapter 2. I expanded this paradigm by adding a reversal learning manipulation in which the others' behavior reversed unannouncedly. Assessing participants' reversal-learning abilities provide insights into how flexibly adolescents can adjust their behavior towards changing levels of others' trustworthiness. Additionally, a non-social version of the task was included to assess whether there are differences between social and non-social trust learning. I aimed to study age-related differences in participants' social and non-social trust learning and trust reversal-learning abilities, and additionally assessed factors underlying individual differences that may affect their social learning abilities.

Chapter 4, builds on the findings from chapter 2 and 3, and describes a study that further examined how adolescents sample information on others' past trust behavior when they have to decide whether to trust them. In a trust sampling paradigm, participants could sample information about the trustworthiness of peers, ranging from always trustworthy to always untrustworthy, before they decide to trust or not (*N* = 157, 10-24 years, same sample as in chapter 3). Using computational modeling, I examined age-related differences in how adolescents used this information to update their beliefs about others' trustworthiness. Together, this setup elaborates on the underlying cognitive processes involved in trusting behavior.

In **chapter 5**, I describe an fMRI study in adolescents (*N* = 74, 9-21 years) in which I examined another type of social learning: prosocial learning. Here, I aimed to investigate how adolescents learn to obtain positive outcomes for self versus for others. By combining computational modeling with functional neuroimaging, I was able to assess in prosocial learning signals in the brain. With this study, I thus aimed to examine age-related and individual differences in prosocial learning on both a behavioral and neural level.

Finally, in **chapter 6** I summarize the results of the empirical studies in this thesis, and provide an overall discussion of the findings and its implications.