Using Novel Methods to Examine the Role of Mimicry in Trust and Rapport

Joanna Hale

Thesis submitted to UCL for the degree of Doctor of Philosophy,

December 2016

I, Joanna Hale, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Signed

Date

Abstract

Without realising it, people unconsciously mimic each other's postures, gestures and mannerisms. This 'chameleon effect' is thought to play an important role in creating affiliation, rapport and trust. Existing theories propose that mimicry is used as a social strategy to bond with other members of our social groups. There is strong behavioural and neural evidence for the strategic control of mimicry. However, evidence that mimicry leads to positive social outcomes is less robust. In this thesis, I aimed to rigorously test the prediction that mimicry leads to rapport and trust, using novel virtual reality methods with high experimental control. In the first study, we developed a virtual reality task for measuring implicit trust behaviour in a virtual maze. Across three experiments we demonstrated the suitability of this task over existing economic games for measuring trust towards specific others. In the second and third studies we tested the effects of mimicry from virtual characters whose other social behaviours were tightly controlled. In the second study, we found that virtual mimicry significantly increased rapport and this was not affected by the precise time delay in mimicking. In the third study we found this result was not replicated using a strict, pre-registered design, and the effects of virtual mimicry did not change depending on the ingroup or outgroup status of the mimicker. In the fourth study we went beyond mimicry to explore new ways of modelling coordinated behaviour as it naturally occurs in social interactions. We used high-resolution motion capture to record motion in dyadic conversations and calculated levels of coordination using wavelet analysis. We found a reliable pattern of decoupling as well as coordination in people's head movements. I discuss how the findings of our experiments relate to theories about the social function of mimicry and suggest directions for future research.

Acknowledgements

Firstly, I would like to thank Antonia Hamilton for all of her guidance during my PhD and for my career beyond. Antonia has been incredibly supportive, encouraging and patient in supervising my research and teaching me new ideas and techniques. I am especially grateful for all the skills I have learnt during my time at UCL, and for the many valuable opportunities I have taken to travel, present my work and get involved in other projects. I would not have achieved so much over the last three years without Antonia's encouragement and confidence in me.

I am very grateful to the UK Ministry of Defence for funding my PhD through Dstl. I especially thank Alex Sandham for her help throughout my PhD as our Technical Partner. Without her advice, support and calming influence, this project would not have been completed on time. I also thank Daniel Richardson for his feedback and suggestions for my research.

It has been a great pleasure and privilege to see the Social Neuroscience group grow around me over the past three years. There are now too many members to thank individually, but I extend my thanks to everyone for sharing their ideas, exchanging favours and creating an atmosphere of friendship and collaboration. I especially thank Harry Farmer and Sylvia Pan, who have been great friends and mentors from the start of my time at UCL, and Paul Forbes and Alex Georgescu for their help and friendship during the time shared in our office. I am also indebted to the Research Assistants and students who carried out data collection for studies in this thesis. My thanks to Dom Oliver, Davide Paoletti, Francesco Buccheri, Maddy Payne and Katie Taylor for all of their hard work and dedication in collecting data to an excellent standard.

Finally, I thank Alex Davan Wetton, who was always there to help when no-one else could, and whose support has meant the most.

Table of Contents

Table of Figures

Table of Tables

Chapter 1. Introduction

After his return from the Amazon rainforest, the naturalist William Henry Bates proposed one of the most important principles in evolutionary theory: mimicry. The butterflies he had been observing mimicked the appearance of poisonous species to avoid being eaten by predators. Broadly speaking, mimicry is to be like another in appearance or behaviour. For humans, mimicry of other people's behaviour seems to be a kind of evolutionary Swiss Army knife, with a wide range of adaptive functions proposed. Rather than helping us avoid predators, each of these suggested functions resonates with the adaptive human qualities of innovation and cooperation noted by Darwin (1859). Adam Smith (1759) suggested that mimicry is a way of feeling what others are feeling, making it easier for us to understand one another and cooperate. Thorndike (1898) highlighted the usefulness of mimicry for learning new abilities from others by doing as they do. During the 20th century, mimicry was proposed to underlie empathy (Allport, 1968; Freud, 1921) and provide a communicative tool to signal togetherness (Allport, 1968; Condon & Ogston, 1966; Piaget, 1946). In light of the adaptive social functions attributed to mimicry, many have gone so far as to claim that our ability to imitate is what makes us human. Over the last two decades, mimicry research in social psychology has referred to this remarkable aspect of human behaviour as the 'chameleon effect'.

The chameleon effect typically refers to unconscious behavioural mimicry, which occurs when one person unintentionally and effortlessly copies another person's posture or body movements without either one being aware (Chartrand & Lakin, 2013; Chartrand & van Baaren, 2009). Mimicry has also been described as 'behaviour matching' (Bernieri & Rosenthal, 1991; Chartrand & Bargh, 1999), and may extend to the contagion of facial expressions (Bavelas, Black, Lemery, & Mullett, 1987; Bavelas, Black, Lemery, & Mullett, 1986; Dimberg, Thunberg, & Elmehed, 2000), moods and emotions (Hsee, Hatfield, Carlson, & Chemtob, 1990; Neumann & Strack, 2000) and characteristics of speech

(Giles & Powesland, 1975; Neumann & Strack, 2000). However, I will use the term 'mimicry' to refer more strictly to the unconscious imitation of body posture and movement.

Mimicry is one of many ways we can coordinate our behaviour with other people during social interactions (Table 1-1). The umbrella term *interpersonal coordination* covers a range of coordinated actions between two people, which can be linked in both space and time. If the actions occur at the same time, they are described as *entrained* or *synchronous* (see Table 1-1, column 1). If one action occurs after a delay but is contingent on the other, this is termed *mimicry* or *imitation* if the actions have the same form, and *complementary* if the form is different (see column 2). We make a distinction between mimicry, which is unconscious and spontaneous, and imitation, which may be deliberate and goal-directed (Bekkering & Prinz, 2002; Bekkering, Wohlschlager, & Gattis, 2000; but see Whiten, McGuigan, Marshall-Pescini, & Hopper, 2009). While the terms above have traditionally occupied separate research literatures, different forms of interpersonal coordination may occur at the same time or dynamically during a real life social interactions. Therefore, while this PhD project specifically focuses on mimicry, I will also draw from literature on imitation, contingent behaviour and synchrony.

Although mimicry between real life social partners may happen reciprocally, in research we typically label one person as the mimicker and one person as the mimickee. With the spotlight predominantly on the mimicker, recent research has built up a large body of evidence about the social and cognitive processes involved in mimicking another person. Data from many sources shows that people tend to spontaneously copy each other (Chartrand, Maddux, & Lakin, 2005; van Baaren, Janssen, Chartrand, & Dijksterhuis, 2009), and the production of mimicry may be modulated by the motivation to affiliate (Chartrand et al., 2005; Lakin, Jefferis, Cheng, & Chartrand, 2003). Neurally, we know that producing mimicry engages inferior parietal cortex and premotor cortex (Grèzes & Decety, 2001; Iacoboni et al., 1999; Molenberghs, Cunnington, & Mattingley, 2009), which are commonly referred to as the mirror neuron system. Top-down control of mimicry in response to the social context has also been linked to neural activation in prefrontal cortex (Brass, Ruby, & Spengler, 2009; Spengler, von Cramon, & Brass, 2010; Spengler, von Cramon, & Brass, 2009; Wang & Hamilton, 2012); As a result, several detailed neurocognitive models of mimicry have been developed which link together processes of mimicry production and top-down control (Brass et al., 2009; Cross, Torrisi, Reynolds Losin, & Iacoboni, 2013; Spengler et al., 2009; Wang & Hamilton, 2012).

Table 1-1. Definitions.

Interpersonal coordination
The degree to which the behaviours in an interaction are nonrandom, patterned, or synchronised in both timing and form (Bernieri & Rosenthal, 1991)

	Synchrony in timing	**Delay in timing**
	Entrainment The behaviour of two moving actors A1 and A2 becomes coupled because they mutually affect each other's behaviour (Knoblich & Sebanz, 2008).	***Contingency*** The extent to which activation of one representation predicts activation of another (Cook, Press, Dickinson, & Heyes, 2010).
Same Form	***Perfect synchrony*** The matching of behaviour in both form and time (Miles, Griffiths, Richardson, & Macrae, 2010), e.g. marching in parade.	***Imitation*** Copying the form of an action (Whiten et al., 2009). Imitation is volitional (Kinsbourne & Helt, 2011) and goal-directed (Bekkering et al., 2000). ***Mimicry*** The automatic imitation of gestures, postures, mannerisms, and other motor movements (Chartrand & Bargh, 1999). Mimicry is not goal-directed (Hamilton, 2008).
Different Form	***General synchrony*** The matching of different behaviours at the same time, e.g. playing of an orchestra.	Complementary actions and other non-matching contingent behaviours, e.g. taking an object from someone’s hand.

In contrast to this detailed evidence about the production of mimicry, we know less about how mimickees perceive and respond to being mimicked. From a theoretical point of view, it is widely assumed that there is a bidirectional link between mimicry and affiliation, such that being mimicked leads to more liking, affiliation and rapport (Chartrand et al., 2005; Chartrand & van Baaren, 2009; Lakin & Chartrand, 2003). This has led to the dominant theory that mimicry has evolved to act as ‘social glue’ (Dijksterhuis, 2005; Lakin et al., 2003) that helps us to bond with members of our social groups by creating smooth, harmonious social interactions (Lakin et al., 2003). Much of the support for this theory comes from evidence that people increase mimicry when they interact with in-group members (Bourgeois & Hess, 2008; Yabar, Johnston, Miles, & Peace, 2006) or otherwise have a goal to affiliate (Cheng & Chartrand, 2003; Karremans & Verwijmeren, 2008; Stel et al., 2010). However, if mimicry truly serves an adaptive social function, then we should also look to mimickees for critical evidence as to whether mimicry creates positive social effects or not.

The aim of this thesis is to rigorously test the claim that being mimicked leads to rapport and trust towards the mimicker. In the following section, I review three different theoretical views about the social function of mimicry. Next, I critically review existing empirical evidence about how people respond to being mimicked outside of awareness. In particular, I will re-examine whether being mimicked consistently leads to positive outcomes such as liking and trust. Finally, I will discuss methodological challenges traditionally associated with studying mimicry effects, and alternative methods which may overcome these challenges. Such methods include programming avatars to mimic in virtual reality and motion tracking participants who have been primed to mimic each other. These methods form the basis for the experiments in this PhD project.

1.1 Theories of Mimicry

1.1.1 Mimicry is Innately Adaptive.

One set of theories suggests that mimicry has evolved as an innate ability with adaptive advantages for communication and bonding. According to this view, mimicry is a matter of nature rather than nurture. Meltzoff and Moore (1977, 1983, 1997) claimed that newborn infants are born with the ability and tendency to imitate adults’ facial expressions. They proposed that this ability comes from a specialised cognitive mechanism for mapping observed actions onto actions that could be performed by the infant's body, called 'active intermodal mapping' (Meltzoff & Moore, 1997). According to a nativist view, this innate mechanism would have evolved through natural selection because mimicry is an important survival tool that provides a basis for social learning and understanding (Meltzoff, 2007a; Meltzoff & Moore, 1997). Most evidence for this account comes from studies of newborn infants and monkeys. For example, Meltzoff and Moore (1977) famously reported that neonates mimicked the facial expressions of adults mere hours after birth. They argued that if such young infants could mimic a range of facial expressions, there must be an innate ability to mimic that could not be explained as a simple reflex. However, subsequent meta-analysis has shown that neonate imitation cannot be reliably replicated for most expressions (Anisfeld, 1996). One exception was tongue protrusion, although a subsequent review also found this effect to be unreliable (Jones, 2009). A recent longitudinal study following 106 infants for nine weeks found no evidence for neonatal imitation, and the data suggested that previous positive findings could be artifacts associated with small sample sizes and limited control conditions (Oostenbroek et al., 2016). Overall, there is little robust evidence that mimicry is an innate ability in newborns and therefore there is not strong support for nativist accounts of mimicry.

Focusing on adult mimicry, another theory proposes that our unconscious tendency to mimic other people is an 'honest signal' (Pentland, 2010). Pentland defines an honest signal as behaviour that is processed unconsciously or is otherwise uncontrollable. As well as mimicry, Pentland suggests that other honest signals include imitation in speech (which he terms 'influence'), amount of overall movement (termed 'activity') and variability in movements and speech ('consistency'). These unconscious patterns of behaviour are assumed to have a biological basis, and it is proposed that they evolved from primate social signalling. According to Pentland (Arena, Pentland, & Price, 2010; Pentland, 2010), honest signals have strong adaptive value because they unintentionally reveal our true attitudes towards people we interact with, particularly in-group versus out-group members. Evidence in favour of the honest signal theory comes from the observation that people mimic unintentionally and outside of their awareness. It is also supported by studies showing that people tend to increase mimicry towards ingroup members (Bourgeois & Hess, 2008; Yabar et al., 2006). However, there is little evidence about how the mimicked party can unconsciously perceive mimicry and recognise it as a signal of liking. In addition, the honest signals theory is directly challenged by evidence that people adjust their levels of mimicry in a goal-directed manner. Following an explicit goal to affiliate, participants increased their level of mimicry outside their own awareness (Lakin & Chartrand, 2003). This finding implies that unconscious mimicry can be affected by deliberate goals, and might therefore be modulated in a Machiavellian way to achieve social influence (Tanner, 2008; Wang & Hamilton, 2012). This evidence therefore undermines the idea that mimicry evolved for honest social signalling.

1.1.2 Mimicry is a Social Strategy.

Another set of theories proposes that mimicry is a strategic communication tool. Lakin et al. (Lakin et al., 2003) describe mimicry as 'social glue' that bonds people together by creating liking, affiliation and rapport (see also Dijksterhuis, 2005). They suggest that the ability to deliberately imitate other people's actions may originally have had adaptive value in helping us to communicate non-verbally. Over time, they propose that this ability might have become automatized and evolved into unconscious behavioural mimicry. According to the social glue theory, mimicry now serves a new adaptive function of creating more harmonious relationships with group members. Thus, the social glue theory makes similar claims to Pentland's honest signals theory. However, the two theories make different claims about the strategic control of mimicry. Whereas Pentland argues that mimicry is adaptive because it is truly honest and uncontrolled, the social glue theory suggests that mimicry is adaptive because people can use it strategically when they need to create affiliation and rapport. This part of the social glue theory is based on two lines of evidence. Firstly, Chartrand and Bargh (1999) showed that people who are mimicked feel rapport towards the mimicker, suggesting that mimicry can be successful in generating rapport. Secondly, Lakin et al.'s (2003) finding that people are more likely to mimic others when they have an explicit goal to affiliate with them suggests that levels of mimicry can be strategically increased. Other studies have reported consistent data on the link from mimicry to liking (Bailenson & Yee, 2005; Kot & Kulesza, 2016; Kouzakova, van Baaren, & van Knippenberg, 2010) and the modulation of mimicry towards people who are liked versus disliked (Stel et al., 2010). However, this literature also contains negative results (Drury & van Swol, 2005; Maddux, Mullen, & Galinsky, 2008; Verberne, Ham, Ponnada, & Midden, 2013). I will go on to give an in-depth critical review of the downstream effects of mimicry in the second section of this chapter.

In line with the social glue theory, the Social Top-Down Response Modulation (STORM) model of mimicry outlines how mimicry can be strategically controlled depending on the social context (Wang & Hamilton, 2012). Rather than giving a theory about the purpose of mimicry, the STORM model outlines how the strategic control of mimicry may be neurally implemented. Specifically, it suggests mimicry is implemented in the mirror neuron system (MNS) and is strategically controlled by medial prefrontal cortex (mPFC). The model is based on evidence from a stimulus-response compatibility (SRC) task. In this task, the participant must close or open their hand while a model in a video makes a compatible or incompatible action. Typically, participants show an imitative effect in which they are faster to respond when the actions are compatible (Brass et al., 2009; Wang, Newport, & Hamilton, 2011). Wang, Newport and Hamilton (2011) showed that this imitative tendency was enhanced when the model in the video maintained direct eye contact, instead of looking away or at her hand. This suggests that levels of mimicry are controlled depending on the presence of direct gaze as a social cue. In addition, Wang, Ramsey and Hamilton (2011) showed that performing the SRC task activated two mirror regions, the STS and IFG, while the observation of eye contact engaged mPFC. The connectivity between the MNS and mPFC increased when participants imitated a model making eye contact, suggesting that the mPFC exerts top-down control over mimicry. Thus, the STORM model gives a detailed account of how mimicry is controlled by top-down mechanisms. Recently, Stel, van Dijk and van Baaren (2016) have proposed a similar cognitive model called Associated Reaction to Action in Context (ARAC). This model also proposes that the automatic elicitation of mimicry is controlled depending on the social context. However, the authors add that this top-down control could depend on previous experiences of reward from mimicking in different contexts, and that it may be limited by cognitive capacity. So far these additions have not been directly tested.

1.1.3 Mimicry is a By-Product of Learned Associations.

In contrast to the theories above, an alternative account claims that mimicry does not have a social purpose but is instead a by-product of learned associations. The associative sequence learning (ASL) theory suggests that mimicry develops through repeated contingent experiences of performing an action and seeing someone else perform the same action (Heyes, 2001; Heyes & Ray, 2000). For example, when a baby waves her arm and sees her own arm moving, this strengthens the neural connectivity between the baby's visual and motor representations of waving an arm. Similarly, when the baby smiles in front of a parent or a mirror and sees a smile back, this strengthens her connection between visual and motor representations of smiling. The ASL theory suggests that, over time, repeated contingent sensorimotor experiences create neural associations between perceiving and doing actions, such that perceiving an action leads to activation of motor neurons that produce the action. This may sometimes produce mimicry (and mirror neurons), but critically the ASL theory also predicts that other contingencies are learned in the same way. Therefore, according to this view, mimicry is simply an evolutionary by-product of associative learning mechanisms. Heyes (2010) argues that mimicry may still have social benefits, but it is not a specialised ability that was selected for a particular social function.

Support for the ASL theory comes from SRC studies showing that automatic imitative behaviour can be reversed or abolished through sensorimotor training (Ray & Heyes, 2011). For example, Heyes, Bird, Johnson and Haggard (2005) gave participants a task where they had to respond with the opposite action to a hand opening or closing. Following relatively brief training on this task over 432 trials, participants no longer showed an automatic imitation response to seeing a hand opening or closing. In another study, participants who were trained to perform foot actions in response to hand actions (and vice versa) showed a similar reduction in automatic imitation following the training (Gillmeister, Catmur, Liepelt, Brass, & Heyes, 2008). Catmur, Walsh and Heyes (2007) showed that imitative tendencies can even be reversed through sensorimotor training of incompatible actions, demonstrating that the involuntary tendency to mimic finger movements can become an automatic tendency to counter-imitate. This research suggests that our automatic tendency to imitate others is malleable and can be changed through relatively brief sensorimotor training. Furthermore, neuroimaging has shown that the reversal of imitative behaviour is reflected in MNS activation (Catmur et al., 2008). Usually, the MNS shows greater activation when hand movements are observed, compared to foot movements. However, after participants were trained to respond with a foot movement when they saw a hand movement and vice versa, the MNS showed more activation when observing foot movements than hand movements. This strongly supports the ASL theory by showing that learning new sensorimotor associations is accompanied by a change in the MNS, suggesting that its mirror properties are formed and changed through learned experience.

1.1.4 Summary and Theoretical Predictions

Current theories about the social function of mimicry fall into three main camps. The first suggests that mimicry is an innate behaviour that has evolved because it is socially adaptive. Nativist theories suggest that mimicry may serve functions associated with social learning and understanding others' minds (Meltzoff, 2007a; Meltzoff & Moore, 1997). The 'honest signals' theory suggests mimicry is adaptive because it provides a true indicator of social attitudes (Arena et al., 2010; Pentland, 2010). Each of these theories has only received weak empirical support. The second theoretical camp suggests that mimicry is used as a social strategy. The social glue theory assumes that mimicry originally evolved to facilitate communication but now serves the function of creating affiliation and rapport (Dijksterhuis, 2005; Lakin et al., 2003). According the social glue theory, people can adjust levels of mimicry strategically. In line with this view, the STORM model provides a detailed neurocognitive account of the top-down control of mimicry depending on social context (Wang & Hamilton, 2012). However, while there is strong behavioural and neural evidence that people use mimicry strategically, it is less clear whether mimicry actually has strategic benefits for social interactions. This is a key limitation, because the third theoretical camp opposes the view that mimicry has a social function. Proponents of this view suggest that mimicry is an evolutionary by-product of learned associations (Heyes, 2001, 2010). The ASL theory proposes that domain-general learning processes lead to visuomotor associations between perceived and performed actions. In some contexts these may give rise to mimicry, but the ASL theory suggests that mimicry itself has no special social purpose.

For critical evidence about whether mimicry does or does not serve a social function, we can look at the effects of mimicry on the mimickee. If being mimicked leads to positive feelings towards the mimicker, this would suggest that mimicry is a useful social strategy and would support the social glue theory. However, if mimicry itself cannot generate any positive social effects then this would undermine its utility as a social strategy, consistent with the view that mimicry is an evolutionary by-product with no social purpose. It is the aim of this thesis to rigorously test whether mimicry leads to positive social effects, with a focus on rapport and trust. In the following sections of the introduction I critically examine the current literature on mimicry effects and highlight methodological issues associated with current approaches, as well as new methods which will be used in this thesis.

1.2 How Do People Respond to Being Mimicked?

In this section, I critically review the current literature on the downstream effects of being mimicked on the mimickee. First, I review the effects of being mimicked on a range of positive social outcomes that have been investigated, including affiliation, trust, prosocial behaviour, self-related feelings and changing opinions. While many studies show positive effects of being mimicked on these outcomes, these results may be tempered by other research showing that the effects of being mimicked are modulated by social characteristics of the mimicker and mimickee.

1.2.1 Positive Responses to Mimicry

Liking. There is a strong consensus that people respond positively to being mimicked. Initially, researchers observed that mimicry during clinical therapy sessions (Cappella & Planalp, 1981; Scheflen, 1964, 1972) and classroom interactions (Bernieri, 1988; Bernieri & Rosenthal, 1991; LaFrance, 1979; LaFrance & Broadbent, 1976) was correlated with reported affiliation, empathy and rapport. Several early experiments manipulated posture congruency and found that confederates who mirrored the posture of participants were evaluated as more similar (Dabbs Jr., 1969; Navarre, 1982), empathic (Maurer & Tindall, 1983) and sociable (Navarre, 1982). Then, in a seminal study, Chartrand & Bargh (1999, Experiment 2) trained confederates to manipulate the level of mimicry in an interaction. Each participant spent fifteen minutes with a confederate, taking turns to describe various photographs. In the mimicry condition, the confederate mirrored participants' posture, gestures and mannerisms; in the control condition, the confederate maintained a neutral posture. At the end of the session, participants who were mimicked rated the confederate as significantly more likeable and the overall interaction as significantly more smooth than participants in the control condition. Following this study, the confederate paradigm became a popular method for studying mimicry effects (Stel, Rispens, Leliveld, & Lokhorst, 2011; Van Baaren & Chartrand, 2005; van Baaren, Holland, Kawakami, & van Knippenberg, 2004), and researchers have worked under the assumption that one of the fundamental effects of mimicry is to increase liking towards the mimicker (e.g. Chartrand & Lakin, 2013; Lakin et al., 2003; Stel et al., 2010).

However, this basic link from mimicry to liking has not been replicated consistently. Eleven studies which measured liking in response to mimicry are summarised in Table 1-2. Five experiments have replicated Chartrand & Bargh's (1999) result using the confederate paradigm (Kot & Kulesza, 2016; Kouzakova, Karremans, van Baaren, & van Knippenberg, 2010; Kouzakova, van Baaren, & van Knippenberg, 2010; Stel et al., 2011, Study 1). One experiment replicated this finding within 'prosocial' but not 'proself' participants (Stel et al., 2011, Study 2). Two experiments using the confederate paradigm failed to replicate the mimicry-liking link, despite reporting significant effects of mimicry on other measures (Drury & van Swol, 2005; van Swol, 2003). Similar results were reported by a much earlier experiment on posture congruency (Dabbs Jr., 1969). Bailenson and Yee (2005) found positive effects of being mimicked using a virtual mimicry paradigm: in their experiment, participants wore a head mounted display (HMD) which let them see a virtual character in an immersive virtual environment. The HMD tracked participants' head movements and the virtual character either mimicked their movement or made head movements recorded from a previous participant, while delivering a persuasive speech. Participants who were mimicked rated the character as more effective on a composite scale which included likability (Bailenson & Yee, 2005), although the weighting of likeability was unclear. Another virtual mimicry study found a positive effect on liking for one out of two virtual characters that mimicked participants in the same way (Verberne et al., 2013). Finally, Maddux, Mullen and Galinsky (2008, Study 2) instructed participants to either mimic or not mimic their partner during a business negotiation task, and found that mimicry did not lead the partners to rate more liking for each other. Overall, only six of 12 studies found a clear mimicry-liking link, and our list does not include studies which have not been published due to negative results. Even the studies which have found positive results report small effect sizes (eta squared close to 0.1).

Trust. The effects of mimicry on trust towards the mimicker appear to be similarly inconsistent. In the same business negotiation task, Maddux et al. (2008, Study 2) found that the amount of time participants self-reported mimicking their partner was significantly correlated with the partner's rating of trust towards the mimicker, and the partner's trust mediated a positive effect of mimicry on the likelihood of negotiating a successful deal. In line with these findings, Verberne et al. (2013) found people rated more trust towards a virtual character that mimicked them, and mimicry also increased participants' willingness to trust the virtual character in an economic investment game. However, they could not replicate these results with a second character and a different decision-making task. This suggests the effects of mimicry on implicit trust behaviour may be mimicker- or task-dependent (Hasler, Hirschberger, Shani-Sherman, & Friedman, 2014). It is worth noting that very few experiments have tested the effects of mimicry on trust, so the mixed results reported here make it hard to draw any firm conclusions. Self-report ratings (e.g. Maddux et al., 2008) and economic trust games where the amount of money or other goods exchanged provides the measure of trust (Verberne et al., 2013) are also relatively explicit measures that require the participant to put a number on how much they trust the mimicker or non-mimicker, and may not capture more implicit aspects of trust. In addition, the economic trust game is highly sensitive to individual differences (Ben-Ner & Halldorsson, 2010; Johnson & Mislin, 2011), which may reduce its sensitivity to differences between a mimicker and non- mimicker. Therefore, mimicry research could benefit from more implicit behavioural measures of trust.

At this point it would also be helpful to discuss how trust and trustworthiness are defined. There is a conceptual distinction between a participant's *trust* towards a target, and the target's *trustworthiness*, although sometimes a participant's ratings of the target's trustworthiness are used to infer their trust (Lyon, Möllering, & Saunders, 2012). There are multiple definitions of trust: a widely-accepted definition in behavioural economics is that trust involves a voluntary transfer of goods or services in a situation where reciprocation is expected but not guaranteed (e.g. Gunnthorsdottir, McCabe, & Smith, 2002), such as the investment game. Under this definition, trust involves risking exploitation by another person, and that person is trustworthy if they do not exploit. In research using virtual characters, and the field of human-computer interaction more generally, trust and trustworthiness are often construed in terms of *cooperation.* Economic games involving a risk of exploitation have been used to measure a user's cooperation with an agent and/or manipulate the agent's cooperativeness (e.g. de Melo, Carnevale, & Gratch, 2011; de Melo, Gratch, Carnevale, & Read, 2012; Krumhuber et al., 2007). In these contexts, cooperation is interpreted as trusting the agent, and has been shown to correlate with ratings of perceived trustworthiness (de Melo, Carnevale, & Gratch, 2013; Krumhuber et al., 2007).

However, in this thesis we adopt a looser cross-disciplinary definition, where trust involves putting oneself in a vulnerable position with positive expectations of another person's intentions or behaviour (Lyon et al., 2012; Rousseau, Sitkin, Burt, & Camerer, 1998). Under this definition, trust does not necessarily involve a risk of exploitation, but could involve a risk of being disadvantaged if the other person is uncooperative, untruthful, unreliable or incompetent. This broad view is able to accommodate nuances in the lay meanings of 'trust' in everyday language. However, it is unclear whether this definition is compatible with the popular model of social cognition proposed by Fiske et al. (Fiske, 2008; Fiske, Cuddy, Glick, & Xu, 2002; Fiske, Cuddy, & Glick, 2007; Fiske, Xu, Cuddy, & Glick, 1999). According to this model, our perceptions of others are organised along two universal dimensions of warmth and competence. Fiske et al. (2007) argue that warmth is closely linked to trustworthiness, although they acknowledge that this link is debated. In contrast, they claim it is widely accepted that competence comprises traits such as efficiency and intelligence. Under this view, perceived trustworthiness may be independent of perceived reliability or competence, although it is less clear where traits like cooperativeness or truthfulness might fit in.

Table 1-2. Studies measuring the effect of mimicry on liking.

Reference	Experimental design	Social interaction	Mimicry condition	Control condition	Mimicry duration	Measure of liking	Participants (*N*)	Reported effect size	Significance (*p*)
Bailenson & Yee (2005)	Between-participants	Persuasive speech	Virtual character mirrored participant head movement	Virtual character displayed previous participant head movement	195 seconds	Agent impression: 13 items[a], including 9-point scale *Not at all likeable – Very likeable*	69	Not reported	<.001
Drury & van Swol (2005)	Between-participants	Debate	Confederate mirrored participant body posture and movement	Confederate moved naturally while avoiding movement related to the participant's movement	10 minutes	7-point scale *Not likeable -Likeable*	78	η^2= 0.01	.90
Kot & Kulesza (2016)	Between-participants	Describing photos	Confederate mirrored participant gestures and mannerisms	Confederate maintained still neutral position	10 minutes	7-point scale *Strongly disagree – Strongly agree (I like the other person / I think she is a nice person / I think she is a good person)*	42	Cohen's d = 0.74	.02
Kouzakova, Karremans et al. (2010)	Between-participants	Mundane tasks, e.g. describing photos and naming depicted animals	Confederate mirrored participant body posture and movement	Confederate moved naturally while avoiding movement synchronous with the participant's movement	5 minutes	7-point scale (*Likeability*)	69	η^2 = .11	.03
Kouzakova Karremans et al. (2010)	Between-participants	Mundane tasks	Confederate mirrored participant body posture and movement	Confederate moved naturally while avoiding movement synchronous with the participant's movement	5 minutes	7-point scale (*Likeability*)	40	η^2 = .12	.03
Kouzakova, van Baaren et al. (2010)	Between-participants	Mundane tasks	Confederate mirrored participant body posture and movement	Confederate moved naturally while avoiding movement synchronous with the participant's movement	10 minutes	7-point scale *Not at all – Very much* (*Likeable*)	72	η^2 = .12	.004

Maddux et al. (2008)	Between-participants	Negotiation	Other participant instructed to mimic participant movements	Other participant not instructed to mimic	45 minutes	5-point scale *How much did you like negotiating with the other person?* *Not at all – very much*	62	Not reported	> .23
Stel et al. (2011)	Between-participants	Description of film fragment	Confederate mimicked participant body posture and movement	Confederate avoided mimicry while keeping other behaviour constant	3 minutes	7-point scale *Did you like your interaction partner? Did you get along with your interaction partner?*	88	$\eta^2 = 0.10$	.01
Stel et al. (2011)	Between-participants	Giving transport directions	Confederate mimicked participant body posture and movement and vocal and facial expressions	Confederate avoided mimicry while keeping other behaviour constant	43 seconds (average)	7-point scale *Did you like your interaction partner? Did you get along with your interaction partner?*	Proself: 22 Prosocial: 27	Proself: Not reported Prosocial: $\eta^2 = 0.08$	Proself: n.s. Prosocial: .05
van Swol (2003)	Within-participants	Debate	Confederate mirrored participant body posture and movement from waist up	Confederate moved naturally while avoiding movement related to the participant's movement	10- 12 minutes	7-point scale *Not likeable -Likeable*	54	Cohen's d = 0.62	.64
Verberne et al. (2013)	Between-participants	Task instructions	Virtual character mirrored participant head movement	Virtual character displayed previous participant head movement	Trial 1: 102 seconds Trial 2: Not reported	Liking: 13 items[a], including 7-point scale *Totally disagree – Totally agree (Likeable)*	40	Trial 1: Not reported Trial 2: $\eta_p^2 = .13$	Trial 1: > .131 Trial 2: .027

[a]Original items can be found in Guadagno & Cialdini (2002)

We return to discuss the relationship between trust, warmth and competence in Chapter 2.

Prosocial and self-related changes. A reliable positive consequence of mimicry is an increase in prosocial behaviour. Following mimicry, participants are not only more likely to agree with an explicit request for help (Guéguen, Martin, & Meineri, 2011), they are also more spontaneously helpful: van Baaren et al. (2004, experiment 1) found that people who were mimicked by an experimenter while taking turns to describe advertisements were more likely to pick up some pens she dropped after the end of the task. In a follow-up experiment, people who were mimicked were also more likely to help an unrelated experimenter (van Baaren et al., 2004). Similar responses were recently demonstrated in infants aged 18 months using an adaptation of the same paradigm (Carpenter, Uebel, & Tomasello, 2013). In other contexts, being mimicked made participants more willing to help an unknown researcher by filling out a tedious questionnaire (Ashton-James, van Baaren, Chartrand, Decety, & Karremans, 2007) and made people passing along a street more compliant with a stranger's request for help (Fischer-Lokou, Martin, Guéguen, & Lamy, 2011). As well as helping, mimicry leads people to donate more money to charity, regardless of whether the charity is connected to the mimicker (van Baaren et al., 2004). People may even be more inclined to vote for prosocial left-wing political parties following mimicry (Stel & Harinck, 2011). Taken together, these findings suggest that mimicry elicits prosocial responses which extend beyond the mimicry interaction (Van Baaren & Chartrand, 2005; van Baaren et al., 2004).

Mimicry also appears to influence or affect the self-construal of the person being mimicked. When completing a 'twenty statements' measure of self-construal (Kuhn & McPartland, 1954), in which people may define themselves by relationships with other people (interdependently) or without reference to others (independently), people reliably provide more interdependent statements following mimicry (Redeker, Stel, & Mastop, 2011; Stel & Harinck, 2011; Stel et al., 2011). Participants who were mimicked also felt closer to others when completing an 'inclusion of other in the self' (IOS) scale (Aron,

Aron, & Smollan, 1992), which depicts increasingly overlapping circles representing self and other (Ashton-James et al., 2007, Experiment 2). As well as feeling closer to others, participants who have been mimicked are more likely to connect objects with their surrounding context and see similarities between photographs which are not systematically related (van Baaren, Janssen, et al., 2009). They also show less divergent thinking and more convergent thinking, which is thought to facilitate collaboration (Ashton-James & Chartrand, 2009). Together, these studies suggest that being mimicked leads to both an interdependent self-construal and prosocial behaviour. Notably, these effects have been demonstrated together (Ashton-James et al., 2007; Catmur & Heyes, 2013; Stel & Harinck, 2011) and Ashton-James et al. (2007, Study 4) found that self-construal mediated the effect of mimicry on prosocial behaviour.

Changing opinions. Being mimicked can change people's opinions and behaviour in a number of ways. Mimicry increases perceived smoothness in an interaction (Chartrand & Bargh, 1999). In addition, people are more likely to disclose intimate information (Guéguen, Martin, Meineri, & Simon, 2013) or give honest answers (Guéguen, 2013) to a confederate who mimicks them. Mimickers are also rated as being more persuasive than non-mimickers (Bailenson & Yee, 2005; Drury & van Swol, 2005; van Swol, 2003), and may sometimes be more successful in swaying people to agree with their opinion (Bailenson & Yee, 2005, but see van Swol, 2003) or to consume and purchase goods (Herrmann, Rossberg, Huber, Landwehr, & Henkel, 2011; Jacob, Guéguen, Martin, & Boulbry, 2011; Tanner, Ferraro, Chartrand, Bettman, & Baaren, 2008). Furthermore, mimicry can improve negotiation outcomes (Maddux et al., 2008): participants who negotiated for around 30 minutes had better personal and joint outcomes when one member of the dyad was instructed to mimic (Maddux et al., 2008). These results suggest that mimicry could indeed be a beneficial social strategy for inducing compliance (Lakin et al., 2003). However, it is possible that confederate mimickers might subtly alter other aspects of their behaviour as well as mimicry; I will discuss this challenge further in the next section.

In addition, increasing conformity of opinions is not always positive. Mimicry can make participants conform to stereotypes consistent with group stereotypes even when those are negative towards the participant and the participant does not endorse them (Leander, Chartrand, & Wood, 2011). Together, these studies suggest that being mimicked may make participants more conformist or likely to agree, with both the good and bad consequences that can bring.

1.2.2 Factors Modulating Positive Responses to Mimicry

Mimicker factors. A large number of factors can alter the general picture that mimicry has positive and prosocial effects. This is particularly clear in situations where people interact with a member of their social outgroup. People typically produce less mimicry towards others who they initially dislike (Stel, Blascovich, et al., 2010), outgroup members (Bourgeois & Hess, 2008; Yabar, Johnston, Miles, & Peace, 2006), and others from a different race (Johnston, 2002). Being mimicked by someone from an outgroup does not seem to have the same prosocial consequences as ingroup mimicry. For example, following mimicry from an ingroup (White) or outgroup (Black) confederates, Dalton et al. (2010, Experiment 2) gave participants a Stroop task as a measure of cognitive resource depletion. The results showed a significant interaction between mimicry and race: participants who were mimicked by a confederate of the same race showed less resource depletion than people who were not mimicked; on the other hand, participants who were mimicked by someone of a different race showed more resource depletion than people who were not mimicked (Dalton et al., 2010). Mimicry by an outgroup member also leads participants to report a room as colder than mimicry from an ingroup member (Leander, Chartrand, & Bargh, 2012, Experiment 3).

Similar effects are found when social status and affiliation is manipulated. Dalton et al. (2010) manipulated status by assigning participants to the role of leader or follower and a confederate to the other role. Participants who were mimicked by a leader showed more resource depletion in a later Stroop task, compared to those mimicked by a follower (Dalton et al., 2010, Experiment 3). Participants who were mimicked by a confederate expressing affiliation showed positive consequences of mimicry, whereas those mimicked by a task-focused confederate did not (Leander et al. 2012, Experiment 1). A plausible explanation for all these effects is that mimicry only has positive consequences in contexts where it is expected. If being mimicked is unexpected, because a partner is an outgroup member or of higher status or not interested in affiliating, then participants do not respond in the same way to being mimicked. However, to test this explanation it would be useful to examine the effects of mimicry from ingroup and outgroup members on more direct measures of liking, rapport or trust.

Mimickee factors. The consequences of mimicry may also depend critically on the personality or other features of the participant being mimicked. In particular, people who are highly ‘proself’ rather than ‘prosocial’ may not respond positively to being mimicked. Stel et al. (2011) defined participants as prosocial if they consistently chose to benefit another player in a game, and proself if they played the game competitively or for individual gain. The prosocial participants reacted positively to being mimicked and indicated more liking towards a mimicker than a non-mimicker; however, this effect was absent in proself participants (Stel et al., 2011). Similarly, although mimicry usually causes people to feel more interdependent, people who naturally have a strong independent self-construal could find it uncomfortable to be mimicked. Highly independent people underestimated the room temperature as a result of mimicry; in contrast, highly interdependent people underestimated temperature when they were not mimicked (Leander et al., 2012, Experiment 2). Individual differences in self-construal can reflect differences in cultural background (Markus & Kitayama, 1991), which may modulate responses to mimicry in a similar way. Sanchez-Burks et al. (2009) showed that US Latino participants, whose culture emphasises social harmony, felt anxious when interviewed by a confederate that did not mimic them, whereas this was not observed in US Anglos. Overall, a variety of findings indicate that people who highly value personal gain or feel independent from others may not show the expected positive reactions to being mimicked.

Social anxiety may also prevent some individuals from responding positively to mimicry. People with high social anxiety tend to focus on themselves and feel awkward during conversations (Heerey & Kring, 2007). Therefore it is not surprising that women with high social anxiety mimic others less than non-socially anxious women (Vrijsen, Lange, Becker, & Rinck, 2010). However, Vrijsen Lange, Dotsch, Wigboldus, & Rinck (2010) also found that women with high social anxiety do not respond positively when they are mimicked by someone else. In their study, women listened to two virtual characters give an opinionated speech; one mimicked participants' head movements and the other did not mimic. Socially anxious women evaluated both characters as similarly likable, friendly and convincing, whereas non-socially anxious women evaluated the mimicking character more highly (Vrijsen, Lange, Dotsch, et al., 2010). This suggests that being mimicked may not have prosocial effects in individuals who focus on themselves due to high social anxiety.

Finally, the prosocial effects of mimicry are expected to break down when people become aware they are being mimicked (Ashton-James et al., 2007; Chartrand & Bargh, 1999; Dalton et al., 2010; Guéguen et al., 2013). This is partly because deliberate imitation can be intended and seen as mockery (Nadel, 2002; van Baaren, Decety, Dijksterhuis, van der Leij, & van Leeuwen, 2009). However, very few studies have directly addressed this expectation, as it is common practice to exclude participants who detected mimicry manipulations from analyses (e.g. Bailenson & Yee, 2005; Cheng & Chartrand, 2003; Drury & van Swol, 2005; van Swol, 2003). Bailenson et al. (2008, Experiment 2) explicitly tested how people respond when they detect they are being mimicked. A virtual character mimicked participants' head movements while delivering a persuasive speech in an immersive virtual environment. Eighty per cent of participants detected they were being mimicked; these participants rated the character as significantly less warm and trustworthy compared participants who did not detect mimicry (Bailenson et al., 2008), suggesting that people may only respond positively to mimicry when they are unaware it is happening. Consistent results come from another study in which participants were told at the end of the experiment (or not) that a confederate had been mimicking them (Manusov, 1992). Participants who were made aware of the mimicry suggested negative reasons for it in a follow-up questionnaire, including manipulation, mockery and being annoying.

1.2.3 Summary

In line with previous literature reviews (Chartrand & Dalton, 2009; Chartrand & Lakin, 2013; Chartrand & van Baaren, 2009; van Baaren, Janssen, et al., 2009), our review suggests there are a variety of ways in which people respond positively to mimicry. A range of studies show that mimicry can change people's perception of the mimicker, including judgements of likeability and trust, although these particular effects are not very reliable. Other studies show that mimicry can change a participant's self-construal, leading an increase in prosocial behaviour, and may also increase agreement and conformist behaviour. The positive results from these studies are generally taken as good evidence for the social glue theory of mimicry. However, research I have reviewed here also shows that positive effects of being mimicked are modulated by characteristics of both the mimicker and mimickee. If characteristics of the mimicker make mimicry seem unlikely, including outgroup membership or high status, then participants do not respond positively to mimicry. Participants who are naturally independent or socially anxious also report less positive effects of mimicry. These findings indicate that mimicry cannot always offer a strategic advantage in all social situations. Overall, the studies reviewed in this section suggest that mimicry can sometimes have positive social effects, in line with the social glue theory, but the evidence is not conclusive.

1.3 Methodological Challenges and Alternative Approaches

Limitations of the current literature may be partly attributed to methodological challenges in studying mimicry effects. In this section I will discuss several major challenges associated with existing approaches to manipulating mimicry and measuring people's responses to being mimicked, as well as challenges for carrying out rigorous experiments. I then introduce two alternative methods for studying mimicry which exploit recent technological advances. The first approach is to generate mimicry from virtual interaction partners. The second is to extracting motion data from naïve participants during a mimicry interaction. Both of these methods are used and developed in the experimental chapters of this thesis.

1.3.1 Methodological Challenges

Manipulating mimicry. The first major challenge in testing the consequences of being mimicked is to achieve a well-controlled manipulation of mimicry. Since mimicry normally occurs unconsciously (Chartrand & Lakin, 2013; Chartrand et al., 2005), it is inherently difficult to generate or eliminate. A compromise is to instruct participants or confederates to mimic in one experimental condition and refrain from mimicking in a control condition. Although this kind of instruction can generate levels of mimicry similar to spontaneous levels (Stel, Dijk, & Olivier, 2009; Stel, van den Heuvel, & Smeets, 2008), this is not guaranteed (Kurzius & Borkenau, 2015). If untrained participants are instructed to mimic it is necessary to perform manipulation checks, such as video recording their behaviour (Stel & Vonk, 2010). Some researchers have relied on asking the participant to report how well they followed the instruction (Maddux et al., 2008), although this check could easily be distorted by demand characteristics or insufficient recall. Even with trained confederates, it may be hard to achieve consistent mimicry or non-mimicry performance (Fox, Arena, & Bailenson, 2009).

It may also be hard to control extraneous variables. The instruction to mimic imposes cognitive demands which could change other aspects of the social interaction, such as emotional understanding (Stel et al., 2009). Furthermore, it is impossible for a confederate to be blind to experimental condition, and hard for them to be blind to the research hypothesis. It is also possible that differences in non-mimicry behaviour from confederates between conditions could influence the experimental results, without confederate or experimenter being aware of this (Doyen, Klein, Pichon, & Cleeremans, 2012; Klein et al., 2012). For example, postural mimicry is normally intertwined with emotional and vocal imitation (Chartrand & Lakin, 2013; Chartrand et al., 2005), and other types of co-ordination like synchrony (Bernieri & Rosenthal, 1991), turn-taking (Pentland, 2010; Wallbott, 1995) and eye contact (Wang, Newport, et al., 2011). Sometimes these behaviours are deliberately included in the mimicry manipulation (e.g. synchrony, Chartrand & Bargh, 1999; facial and vocal imitation, Stel et al., 2011), but researchers wishing to control for these variables must usually video their experiment and code the behaviour *post hoc* (e.g. Chartrand & Bargh, 1999; Drury & van Swol, 2005; Sanchez-Burks et al., 2009; van Swol, 2003). This approach has limitations because it is time consuming and may not be perfectly accurate.

There are also challenges associated with achieving a good control condition. In some paradigms, the control condition is defined as non-mimicry, i.e. neutral movements (e.g. Chartrand & Bargh, 1999; Kouzakova, Karremans, et al., 2010; van Baaren et al., 2004). Although the confederate is allowed to move, this condition usually involves sitting relatively still in a neutral pose. This could result in less than typical amounts of movement, which might be perceived as unnatural compared to the level of movement in the mimicry condition. In other paradigms, the control condition involves anti-mimicry, i.e. deliberately dissimilar movements (e.g. Ashton-James et al., 2007; Hasler, Hirschberger, Shani-Sherman, & Friedman, 2014). This approach may lead to a better match between the amount of movement in mimicry and non-mimicry conditions, compared to non-mimicry paradigms. However, anti-mimicry may be a poor ‘baseline’ because it involves contingent opposite behaviours. Non-mimicry and anti-mimicry conditions have been shown to generate different effects; for example, people bought significantly more products when they were not mimicked compared to anti-mimicked (Kulesza, Szypowska, Jarman, & Dolinski, 2014). Therefore, researchers need to consider the appropriate control condition to use.

Beyond the specific field of mimicry, researchers have criticised the use of confederates as conversational partners in studies of language and social interaction. Kuhlen and Brennan (2013) highlight four concerns: firstly, that confederates may introduce inadvertent bias if they are not blind to the study's hypothesis; secondly, that confederates risk detection as false participants, potentially changing participants' behaviour; thirdly, that confederates may fail to give natural social cues to participants, due to the extra knowledge they have about the interaction context; and finally, that confederates are often required to follow a script, which could result in unnatural utterances. These concerns may be reflected in evidence that conversations with confederates or experimenters produce different social effects compared to free conversations between naïve participants. A recent review of studies which examined how much participants gestured depending on whether or not they could see their conversation partner (Bavelas & Healing, 2013) found that significant differences were reported by seven studies where the partner was a confederate with constrained responses, and no significant differences were reported by seven studies of free conversation between naïve participants. In a direct test of referential gaze patterns, Brown-Schmidt and Tanenhaus (2008) found that participants also fixated differently on items that were mentioned by an experimenter versus another participant, even though participants could not see each other or the experimenter. This suggests that there may be core differences between conversation with a true partner compared to a confederate. Overall, evidence suggests there are substantial limitations to using confederates as interaction partners in experiments aiming to test spontaneous social phenomena such as the production and downstream effects of mimicry.

Modulators of mimicry effects. Another challenge is to test how mimicry effects are modulated by social contexts and characteristics of the mimicker. This challenge particularly applies to confederate paradigms which manipulate mimicry within a live interaction (e.g. Chartrand & Bargh, 1999; Stel et al., 2011; van Baaren et al., 2004), because it is necessary to (a) find the right confederate and (b) train that person to perform appropriately. Confederate features such as race, gender and age may all affect mimicry, but would be hard to control in a traditional research setting. For example, a researcher interested in how age moderates mimicry effects could not employ a child confederate for ethical and practical reasons.

Choosing implicit or explicit measures of mimicry effects. Another major challenge is to find valid ways of measuring how participants respond to being mimicked. In particular, we can choose between explicit and implicit measures. Explicit measures refer to ones where the participant is aware of what is being measured, whereas implicit measures are ones where the participant is either unaware of what is specifically being assessed or is unaware of their own attitude or behaviour (Petty, Fazio, & Brinol, 2012). Note the distinction refers to measures, and the relationship between implicit and explicit attitudes at the cognitive level is a matter of debate (Fazio & Olson, 2003). In mimicry studies, explicit measures usually involve questionnaires or ratings about the mimicry interaction. The advantage of questionnaires is that they are easy to administer and widely used (e.g. Bailenson & Yee, 2005; Chartrand & Bargh, 1999; Stel et al., 2011; Stel & Vonk, 2010; Vrijsen, Lange, Dotsch, et al., 2010). However, there are no standardised rating scales for many of the constructs of interest in mimicry research, such as liking and rapport (see Table 1-2). Note that explicit measures do not have to use standard questionnaire items, and may sometimes involve pictorial scales or analogies. For example, IOS scale (Aron et al., 1992) uses overlapping circles to illustrate interpersonal closeness, and the ‘feeling thermometer’ uses a temperature scale to measure feelings of warmth or liking; however, these measures have not consistently shown positive mimicry effects (Ashton-James et al., 2007; Hasler et al., 2014; Hogeveen, Chartrand, & Obhi, 2014). Self-report questionnaires and ratings have the limitation of being open to bias in interpretation and responses (Ben-Ner & Halldorsson, 2010; Lyon et al., 2012), and are unsuitable for measuring implicit attitudes. This means that explicit responses to questionnaires may correspond poorly to actual behaviour (Armitage & Christian, 2003).

Therefore, it is also useful to measure implicit behavioural responses to mimicry. A wide variety of implicit behavioural measures have been used in previous research, including Stroop task reaction times (Dalton et al., 2010), estimates of room temperature (Leander et al., 2012), the number of pens picked up (van Baaren et al., 2004), and seat choice (Ashton-James et al., 2007). These measures have the advantage of measuring participants' implicit reactions to being mimicked, but are not very closely related to the mimicry itself and could be influenced by other factors. The variety of implicit measures in existing research can also be seen as an advantage and a disadvantage (Daniël Lakens, Schubert, & Paladino, 2016). On the one hand, replicating mimicry effects across different measures can demonstrate robustness and generalizability. On the other hand, it makes it difficult to compare outcomes and possible mechanisms across different studies. Mimicry research would therefore benefit from establishing implicit behavioural measures that can be easily shared and replicated by multiple research groups. Better evaluation of the relationship between implicit and explicit mimicry responses within the same experiment could also provide important insights for understanding cognitive mechanisms involved in responding to mimicry.

Scientific rigour. As well as specific challenges for manipulating mimicry and measuring its effects, a major issue for psychology research is the need for greater rigour in designing and reporting experiments. Here I will briefly outline four interrelated issues affecting mimicry research. Firstly, it is increasingly recognised that traditional experiments in psychology may lack statistical power (Chase & Chase, 1976; Cohen, 1962; Tressoldi, 2012). Our brief review of the mimicry-liking link suggest an approximate average effect size of $\eta^2 = .01$ and an average sample size of 60 participants. Most studies have used a between-subjects design, possibly to reduce participant awareness of the experimental conditions, although this is not necessarily an effective precaution (Lambdin & Shaffer, 2009). A power-analysis (G*Power, version 3.1.7, 2013) suggests that detecting an effect of $\eta^2 = .1$ with a between-groups design would require 120 participants per group. Detecting a similar effect size with a within-subjects design would require only 22 participants (c.f. Cohen, 1992). As new factors are introduced, increasingly large participant samples must be recruited for between-subjects experiments to achieve sufficient experimental power. Between-subjects paradigms are also hard to adapt to fMRI to allow neuroimaging. Therefore, it may not be feasible to study how mimicry effects vary across different contexts and individuals using traditional between-subjects paradigms.

Secondly, traditional experimental psychology approaches focus on null hypothesis significance testing (NHST), which recent critiques have associated with an over-emphasis on *p*-values rather than effect sizes. The problem with focusing on *p*-values is that they may represent a relatively arbitrary level of ‘significance’ and do not provide information about the size of the effect (Lakens, 2013). In general, effect sizes were reported in almost all studies we reviewed (with some exceptions, e.g. see Table 1-2), but very few authors discussed what these mean or whether their research designs were based on previous effect sizes (as is best practice). Another problem with NHST is that it is very hard to establish a null effect of being mimicked. An alternative approach which avoids these limitations of NHST is to carry out Bayesian statistical analysis (Nathoo & Masson, 2016; Wagenmakers, 2007). Bayesian analysis calculates the likelihood of the observed data under the null hypothesis and under the alternative hypothesis and the resulting Bayes factor can be interpreted as a measure of how much evidence the data provides in favour of either hypothesis.

Thirdly, because of the traditional emphasis on *p*-values in psychology research and the publication bias in reporting significant results (Ferguson & Heene, 2012; Francis, 2012), it can be very tempting for researchers to present exploratory analyses as though they were confirmatory. This is particularly a problem if (a) the researcher carried out many statistical tests they do not report and do not correct for before finding one that was significant (‘*p*-hacking’) or (b) the researcher continued recruiting more participants until they obtained a significant result. In both cases, this can lead to the reporting of false-positives. In order to counter this kind of practice, there is growing support for pre-registration of experimental methods (Jonas & Cesario, 2015; van ’t Veer & Giner-Sorolla, 2016). In the most strict form, this involves writing a full specification of the sample, data collection method and planned analyses prior to data collection, and submitting a registered report. Some journals will agree to publish the final results based on the merit of a registered report (e.g. *Cortex*, Elsevier, 2014). However, it is also possible to specify analyses on an existing dataset, or pre-register other aspects of the research design using platforms such as the Open Science Framework.

Finally, another issue in psychological research is a lack of replication (Francis, 2012; Maxwell, Lau, & Howard, 2015). This has particularly caused controversy over social priming effects, in which subtle social cues are thought to unconsciously influence cognition or behaviour. For example, studies have reported that participants felt more or less close to family members after plotting points that were close or further apart (Williams & Bargh, 2008), walked more slowly after reading words related to the elderly (Bargh, Chen, & Burrows, 1996), and donated less money after reading money-related words (Vohs, Mead, & Goode, 2006). However, this area of research recently received strong criticism in an open letter from Daniel Kahneman (Yong, 2012) that highlighted widespread doubt about the validity of classic social priming results. Kahneman argued that wider problems with replication in psychology (including the file drawer problem) could especially affect social priming research because this field has traditionally favoured conceptual replications rather than exact repetitions of the same method. Several direct replications of classic experiments have failed to find significant results (e.g. Doyen et al., 2012; C. R. Harris, Coburn, Rohrer, & Pashler, 2013; Pashler, Coburn, & Harris, 2012; Shanks et al., 2013), contributing to concerns about the original effects. However, John Bargh responded to criticisms (Bargh, 2012) by pointing out that some social priming effects such ‘elderly walking’ have been independently replicated (Cesario, Plaks, & Tory, 2006; Hull, Slone, Meteyer, & Matthews, 2002). Others have pointed out that single replications can lack the power needed to replicate classic effects (Maxwell et al., 2015).

In the mimicry literature that we have reviewed it is often accepted that there are consistent effects across studies. However, as we mentioned above, the wide variety of dependent measures (and manipulations) used across separate experiments can make it difficult to establish the reliability of particular effects (Lakens et al., 2016). In the case of liking, we found that the effect has not been consistently replicated across studies, but these employed different measures (Table 1-2). It is a matter of debate how closely methods should match in order to be considered a replication (Maxwell et al., 2015; Stroebe & Strack, 2014), with some arguing that direct replication is essential for establishing reliability and others arguing that looser replications demonstrate the generalizability of an effect. We would consider both types of replication to be useful, and therefore researchers could contribute to the field by both seeking to replicate the results of other research groups using novel methods and establishing that their own methods give replicable results over multiple experiments.

1.3.2 Alternative Approaches

Virtual mimicry. Generating mimicry in a virtual reality (VR) setting is an alternative method that overcomes many traditional challenges associated with achieving a controlled mimicry manipulation and investigating modulating factors such as the mimicker's appearance. VR involves the computer simulation of a seemingly real environment, and first had applications in flight simulation and medical training (Rheingold, 1991) following the development of 3D computer graphics in the 1960s (Loomis, Blascovich, & Beall, 1999). Two main hardware systems have been designed to display virtual environments to a user (Loomis et al., 1999); head-mounted displays (HMDs) are perhaps the most widely used. HMDs display stereo images via head-mounted eyepieces and track the user's head movements in order to update the view and simulate looking around a 3D space. Until the 1990s, visual realism was limited by slow computer graphics processing, but this has vastly improved so that now even mobile phones can deliver high-resolution HMD displays (e.g. Samsung GearVR and Google Cardboard systems). The other main system is the CAVE (cave automatic virtual environment) which is a room-sized structure made from projector screens with incorporated motion sensors. A user wearing 3D glasses can walk around the CAVE and experience a virtual environment that updates when they move. VR displays are often combined with headphones to provide audio feedback and haptic devices which deliver vibrations or force feedback to the hands or body in order to add to the sensory realism of the virtual environment. However, as well as sensory realism, the virtual environment also needs to seem physically interactive in order to make someone feel like it is reality, which is termed ‘presence’ by computer scientists. Physical interaction can be simulated by sensing a participant’s movement and updating the environment accordingly. For example, a basic form of interactivity is updating the visual display according to head motion sensors; a more sophisticated form involves applying body-worn sensors to the user so that their movement can drive the actions of a virtual avatar which they embody in the virtual space.

Similarly, in order to generate socially realistic virtual humans (termed ‘virtual characters’ or ‘agents’), they need to seem socially interactive. The term ‘co-presence’ is used to describe the feeling of being with another person in a virtual environment and this is typically assessed with questionnaires (Sanchez-Vives & Slater, 2005). Whereas physical interactivity can be simulated by sensing the user’s physical movements, social interactivity is simulated by sensing social cues from the user and triggering socially appropriate responses from a virtual human. Social sensing is more complicated than physical sensing, because we have an incomplete understanding of physiological, bodily, facial and speech cues and what they mean in real life social interactions. Nevertheless, even very basic forms of social interactivity can be enough to generate feelings of social realism. For example, behavioural and fMRI data show that a very basic virtual human head with an expressionless face can be perceived as socially interactive and neurally rewarding if it makes eye movements that are contingent on where the participant is looking (Schilbach et al., 2010, 2011; Wilms et al., 2010).

Virtual reality is a therefore a useful tool for studying social interactions, as it has been shown in a range of settings that people usually react to virtual characters similarly to how they would with real people (Bailenson, Blascovich, Beall, & Loomis, 2001;

Donath, 2007; Garau, Slater, Pertaub, & Razzaque, 2005; Reeves & Nass, 1996). For example, people maintain appropriate social distance from virtual characters (Bailenson et al., 2001; Bailenson, Blascovich, Beall, & Loomis, 2003; McCall & Singer, 2015), imitate their behaviours (Vrijsen, Lange, Becker, et al., 2010) and show distress when they are harmed (Pan, Banakou, & Slater, 2011; Pan & Slater, 2011). In a public speaking context, people reacted anxiously to an audience of visibly bored virtual characters (Slater, Pertaub, & Steed, 1999), and in a replication of the classic Milgram obedience experiment, some participants were unwilling to 'shock' a virtual human (Slater, Antley, et al., 2006). Given these reactions, social psychologists have used virtual characters to study a range of social phenomena, including prosociality (Gillath, McCall, Shaver, & Blascovich, 2008), trust (McCall & Singer, 2015; Verberne, Ham, & Midden, 2015), persuasion (Bailenson & Yee, 2005; Guadagno, Blascovich, Bailenson, & Mccall, 2007; McCall & Blascovich, 2009; Zanbaka, Goolkasian, & Hodges, 2006), embodiment of other people (Banakou, Groten, & Slater, 2013; Peck, Seinfeld, Aglioti, & Slater, 2013; Slater et al., 2009), social biases (Hasler et al., 2014; McCall, Blascovich, Young, & Persky, 2009; Peck et al., 2013; Zanbaka et al., 2006) and social anxiety (Pan, Gillies, Barker, Clark, & Slater, 2012; Vrijsen, Lange, Becker, et al., 2010; Vrijsen, Lange, Dotsch, et al., 2010).

Virtual reality has also been used to study mimicry effects. Bailenson & Yee (2005) first developed a method for virtual mimicry. They fitted participants with an HMD which displayed a virtual character in immersive 3D and a head-mounted sensor which could also track the participant's head rotation. Bailenson & Yee (2005) then programmed the virtual character to deliver a persuasive speech and mirror the participant's head movements like a reflection, with a delay of four seconds between the participant's movement and the character's movement. They suggest that delay was optimal for maximising mimicry responses while minimising detection, based on a previous pilot study with 41 participants (Bailenson, Beall, Loomis, Blascovich, & Turk, 2004). To achieve a control condition where the character did not mimic the participant, head movements recorded from the previous participant while being mimicked were applied to the character instead. So far, this and other virtual mimicry studies have only tracked head movements (Bailenson et al., 2008; Verberne et al., 2013, 2015; Vrijsen, Lange, Dotsch, et al., 2010), but tracking could be extended to the whole body using current sensor technologies.

Virtual mimicry has the advantage of high control over the mimicry manipulation, because virtual characters are 'reverse engineered' to only perform necessary behaviours (Fox et al., 2009), such as speaking, blinking and mimicking or not mimicking (Bailenson & Yee, 2005). The mimicry and control conditions are also well-matched, because the motion of the non-mimicking character is yoked to the movements the previous participant made while they were being mimicked. Furthermore, the mimicry interaction can be perfectly replicated using the same computer code (Verberne et al., 2013), while characteristics of the character and the virtual environment can be endlessly tailored. For example, the researcher who wanted to investigate age could program a child character to mimic participants (Banakou, Groten, & Slater, 2013). Finally, it may be more feasible to measure real-time responses to mimicry using virtual reality. Motion tracking devices and physiological sensors for heart rate, skin conductance and breathing rate have been used to measure how people physically respond during a VR scenario (e.g. Bailenson et al., 2008; McCall, Hildebrandt, Hartmann, Baczkowski, & Singer, 2016). Alternatively, researchers can play back recorded segments of the participant's virtual experience when they make ratings afterwards (McCall, Hildebrandt, Bornemann, & Singer, 2015), or even allow them to rate their experiences in real time using a virtual interface.

However, there are some disadvantages to virtual mimicry. First, virtual mimicry is designed to be an all-or-none behaviour, which cannot easily be ramped up or down within a single interaction in the same way as natural human mimicry. Second, the virtual characters must also be programmed with other aspects of natural social interaction (e.g. joint gaze) to make them socially realistic. This can be technically difficult to implement.

On the other hand, the precise control of every individual social behaviour in virtual reality can be described as an advantage because it allows us to test the impact of each behaviour separately. Even when virtual characters are very limited in their other social behaviours, existing VR studies demonstrate they can achieve mimicry effects similar to human confederates (e.g. Bailenson & Yee, 2005; Verberne et al., 2013; Vrijsen, Lange, Dotsch, et al., 2010).

Motion tracking naïve participants. Another way to avoid problems associated with instructed mimicry is to record mimicry as it spontaneously occurs during interactions between two participants in a laboratory, while neither of them knows that mimicry is under investigation (e.g. Hess & Bourgeois, 2010). In this context, levels of mimicry can been monitored through detailed video scoring by trained coders (Condon & Ogston, 1966; Heerey & Crossley, 2013; Heerey & Kring, 2007; Kendon, 1970; Kurzius & Borkenau, 2015). Typically, this involves multiple coders rating the videotape frame-by-frame for a set of pre-specified behaviours, such as posture shifts or specific facial expressions (Condon & Ogston, 1966; Heerey & Crossley, 2013; Messinger, Fogel, & Laurie, 1999, 2001). Scoring videos in this level of detail has the advantage of generating very rich data about the social interaction without the need for intrusive equipment such as motion tracking sensors. It is also possible for trained coders to rate aspects of behaviour that would be hard to automatically capture with current motion tracking technologies. For example, Messinger et al. (1999) and Heerey et al. (Heerey & Crossley, 2013; Heerey & Kring, 2007) have used trained coders to rate the presence of genuine versus polite smiles, which current facial tracking algorithms would not be able to distinguish. Given that mimicry may involve variable timing and degree of matching, and that people may switch between mimicker and mimickee roles in natural interactions, trained coders might also be more accurate at assessing mimicry compared to automatic software algorithms.

However, detailed video scoring by specially trained coders can be prohibitively time-consuming. The ‘thin slice’ method is an alternative approach that drastically reduces the time taken to score a videotaped interaction and does not require trained coders. A ‘thin slice’ is a short video clip, defined as ‘a brief excerpt of expressive behaviour sampled from the behavioural stream’ (Ambady, Bernieri, & Richeson, 2000, p. 203). For example, this could be a 10-second video clip of a conversation between two people. When participants are asked to rate a thin slice depicting a social interaction, their first impressions are highly predictive of the actual outcomes of the interaction, such as levels of rapport (Ambady & Rosenthal, 1992), jury decisions (Parrott, Brodsky, & Wilson, 2015) and job interview success (Nguyen & Gatica-Perez, 2015). Interestingly, the duration of the clip has little impact on the predictive power of thin slices, which can be demonstrated for clips as short as 6 seconds (Ambady et al., 2000; Ambady & Rosenthal, 1993). Several research groups have therefore used thin slice ratings to assess synchrony and interpersonal coordination in videotapes of dyadic interactions (e.g. Bernieri, 1988; Bernieri, Steven, & Rosenthal, 1988; Ramseyer & Tschacher, 2011). However, the accuracy of thin slices ratings also depends on the level of expressivity of people in the videotape (Ambady et al., 2000) and raters may take into account many different social signals when rating the interaction, even if they are instructed to focus only on one aspect such as synchrony (Cappella, 1981; Lumsden, Miles, & Macrae, 2012).

Motion tracking methods. As an alternative to manual video coding, automatic recording and analysis techniques are now available to assess levels of spontaneous mimicry or interpersonal coordination between two (or more) participants. One option is to video record the participants as they interact, and then use automatic image processing algorithms to extract data about how each participant moved. For example, a participant’s total body movement can be tracked with frame-differencing methods, which look at the change in video pixels from one frame to the next (Paxton & Dale, 2013). As long as the video background is stable, a change in pixels can be attributed to a participant moving. A recent study using this method showed that true interaction partners synchronise with each other more than randomly paired videos of different partners (Fujiwara & Daibo, 2016), replicating findings from two earlier studies which relied on video ratings of synchrony from untrained observers (Bernieri, 1988; Bernieri et al., 1988). Frame-differencing methods have the advantage specialist equipment is not needed, and have been used in several clinical studies of bodily synchrony (Kupper, Ramseyer, Hoffmann, Kalbermatten, & Tschacher, 2010; Nagaoka & Komori, 2008; Ramseyer & Tschacher, 2008, 2011). However, they are poorly suited to studying mimicry of specific actions. Currently there are few established methods for automatically detecting particular body movements from video footage, although this is an area of rapid development (Michelet, Karp, Delaherche, Achard, & Chetouani, 2012; Sun, Nijholt, Truong, & Pantic, 2011; Sun, Truong, Pantic, & Nijholt, 2011).

Another option is to use a motion tracker to directly record the movements of each participant. There is now a very wide range of motion tracking systems available, although three types are commonly used in social interaction research. First, optical systems triangulate the position of a marker on the participant's body using two or more cameras which can 'see' the marker either because it is reflective (passive marker) or because it emits light intermittently (active marker). Optical systems have the disadvantage that markers can easily become obscured from the cameras' view (Poppe, Zee, Heylen, & Taylor, 2013). Second, magnetic systems also use markers on the participant's body, but they detect the position and rotation of the markers within a weak magnetic field (e.g. see Feese, Arnrich, Tröster, Meyer, & Jonas, 2011, 2012). Magnetic systems tend to be more precise than optical systems, require less calibration, and the markers cannot become obscured. Third, 3D-camera systems such as Kinect (Microsoft) use a camera with an infrared depth sensor to capture a 3D image of a participant's body without the need for markers, and additional software is used to interpret the image (e.g. see Won, Bailenson, Stathatos, & Dai, 2014). This type of system is ideal for facial capture (e.g. Li, Mian, Liu, & Krishna, 2013) as it does not require markers to be attached to the face. However, it also requires lengthy calibration to achieve precise readings.

Analysis of motion tracking data. If the motion of two participants is directly recorded in two time series, there are several options to automatically analyse the level of mimicry or interpersonal coordination in their movements. The two time series can be analysed according to the timing, frequency or both timing and frequency of motion (Fujiwara & Daibo, 2016; Grinsted, Moore, & Jevrejeva, 2004; Issartel, Marin, Gaillot, Bardainne, & Cadopi, 2006). The timing can be analysed by performing a cross-correlation and seeing at what time lag the two time series are most highly correlated. This peak correlation would indicate the timing at which the two participants tended to match each other's motion, and is therefore suitable for studying mimicry or behaviour matching (Fujiwara & Daibo, 2016). The frequency of each person's movements can be examined using spectrum analysis. Spectrum analysis decomposes the data from a time series into different constituent frequencies in order to see how much movement there was at each frequency. Fourier analysis is a well-known method of spectrum analysis. It assumes that the time series follows repetitive patterns and stable frequencies over time, and is therefore suitable for studying how much synchrony or entrainment there is between two participants making repetitive movements (Fujiwara & Daibo, 2016; Issartel et al., 2006), such as walking in step or rocking in a rocking chair. Finally, wavelet analysis can be used to examine both the timing and frequency of motion (Grinsted et al., 2004; Issartel et al., 2006). A wavelet is a localised oscillation or 'blip' in time, and can be used to characterise the frequency (oscillation) and timing (localised blip in time) of the movements in a time series. Wavelet analysis therefore expands a time series into time-frequency space, with the advantage that it can be used to track how the frequency spectrum of the original signal changes over time. This makes wavelet analysis a more suitable tool than spectrum analysis for studying coordination in spontaneous and non-repetitive social interactions such as having a natural conversation (Fujiwara & Daibo, 2016; Issartel et al., 2006). Therefore, studies have wavelet analysis to evaluate interpersonal coordination during free conversation (Fujiwara & Daibo, 2016), musical improvisation (Walton, Richardson, Langland-Hassan, & Chemero, 2015), and telling knock-knock jokes (Schmidt, Nie, Franco, & Richardson, 2014).

Priming naïve mimicry. A remaining challenge for studying mimicry in naturalistic conversations is how to manipulate the level of mimicry. If two participants are not instructed about how to behave, then there is no guarantee that they will mimic one another. If there is little or no mimicry, this would make it very difficult to test the downstream effects of natural mimicry on the people interacting. One way to overcome this problem could be to increase or decrease levels of mimicry through subliminal priming. Priming involves the unconscious or unintentional facilitation of a particular behaviour, such as mimicry, through exposure to a particular type of stimulus or event (Molden, 2014).

The possibility of priming mimicry has been demonstrated across several studies using a scrambled-sentences priming task. In one study, participants were given 18 sentences with five words in the incorrect order, and were asked to make a grammatically correct four-word sentence (van Baaren, Maddux, Chartrand, de Bouter, & van Knippenberg, 2003). In one condition, the sentences contained words related to an interdependent self (e.g. ‘cooperate’) and in the other condition, words related to an independent self (e.g. ‘unique’). Participants in the interdependent condition later spent more time mimicking a target behaviour (pen-playing) displayed by a confederate in a separate task. Other studies have tested the effect of scrambled sentence priming on the automatic imitation of finger tapping movements in SRC paradigms. Two studies compared prosocial words with antisocial words and found that the prosocial priming led to greater imitation than antisocial priming (Cook & Bird, 2011; Leighton, Bird, Orsini, & Heyes, 2010), ,consistent with other research that has reported increased mimicry of target behaviours from a confederate following unsuccessful affiliation (Lakin & Chartrand, 2003) or third party ostracism (Lakin, Chartrand, & Arkin, 2008; Over & Carpenter, 2009). Subsequent research showed that the effect of prosocial and antisocial scrambled sentences seems to be reliable but depends on whether the sentences describe first person or third person events (Wang & Hamilton, 2013). Taken together, these studies provide converging evidence from multiple paradigms to suggest that mimicry may be reliably increased by first-person prosocial stimuli or third-person antisocial scrambled sentences.

Whereas some social priming effects have failed to replicate (e.g. Doyen et al., 2012; C. R. Harris et al., 2013; Pashler et al., 2012; Shanks et al., 2013), the effects of prosocial and antisocial primes (specifically, scrambled sentences) appear to be relatively robust. The fact they have been replicated across different research groups measuring different mimicked actions (e.g. foot versus finger movements) suggests the scrambled sentence task could be a reliable and flexible way of manipulating mimicry within participants, with the major advantage that it could allow us to study mimicry as it spontaneously occurs. This is important, because the majority of empirical evidence we have about how people respond to mimicry comes from studies where mimicry was artificially instructed, which may lead to behaviour that diverges from spontaneous mimicry interactions. Priming is also suitable for within-participants designs, as participants can be subtly primed with different conditions that induce different levels of mimicry over the course of one experiment (e.g. Wang & Hamilton, 2013). However, the few initial studies that have demonstrated priming of mimicry all showed effects on specific actions such as foot tapping or finger tapping that were performed by a confederate or stimulus figure (Lakin & Chartrand, 2003; van Baaren et al., 2003). It remains to be tested whether the same results can be conceptually replicated within more naturalistic contexts such as free conversations.

1.3.3 Summary

To make progress in understanding how people respond to mimicry, new research will need to overcome the major traditional challenge of manipulating mimicry in a naturalistic way under controlled lab conditions, as well as other issues highlighted above. Virtual mimicry and motion tracking naïve participants are two modern approaches that show promise in overcoming some of these challenges, with the respective advantages of strong experimental control and high ecological validity. Existing proof-of-principle studies demonstrate the validity of these approaches. Therefore I use and develop them in this thesis to rigorously test the claims put forward in the social glue theory of mimicry, and to investigate new factors which could tell us about the cognitive processes behind people's responses to being mimicked.

1.4 Overview of Experimental Chapters

In this chapter I have reviewed competing theories about the social function of mimicry. On the one hand, the social glue theory of mimicry suggests that mimicry is an adaptive social strategy for creating rapport. On the other hand, the ASL theory suggests that mimicry is simply an evolutionary by-product that has no special social purpose. Although behavioural and neural evidence suggests people produce mimicry strategically, in line with their social goals, to accept the claim that mimicry is an adaptive social strategy we would also need to show that being mimicked creates positive feelings towards the mimicker. The current literature on effects of being mimicked, which I have reviewed in this chapter, provides some support for the social glue theory. However, existing research in this area has methodological limitations. In particular, most studies have used confederate paradigms, in which it is difficult to isolate the effects of mimicry and avoid all other social confounds. Therefore, we have proposed two alternative approaches to studying mimicry which overcome some of the challenges associated with confederate paradigms: (1) programming virtual characters to mimic participants, and (2) motion tracking naïve participants.

The aim of this thesis is to rigorously test the claim that being mimicked leads to rapport and trust towards the mimicker. It is central to the social glue theory that being mimicked should lead to feelings of rapport, and this outcome has been tested in many previous studies. We aim to see if previous effects can be replicated using a strict approach. In contrast, relatively few studies have tested whether mimicry leads to trust, although several have noted the advantages of mimicry for persuasion and compliance.

One challenge for investigating trust outcomes is that we lack behavioural methods for measuring trust towards specific individuals. We aim to address this problem in Chapter 2, by developing a new behavioural task for measuring trust. While Chapter 2 focuses on measuring trust, Chapters 3 and 4 move on to the main focus of this thesis and use a novel virtual mimicry paradigm to test the effects of mimicry on rapport and trust. Finally, in Chapter 5 we widen our scope to investigate the parameters of interpersonal coordination in natural conversations by motion tracking naïve participants. A central theme throughout the thesis is the use and development of novel methods for investigating mimicry.

Specifically, the studies in this thesis address the following questions:

1. How can we measure trust behaviour towards specific targets?

Chapter 2 develops a new behavioural task for measuring trust towards specific targets. In this task, the participant is immersed in a 3D virtual maze. They must find the way out of the maze and can choose whether to ask for and follow advice from different virtual characters. Across two experiments we show that the virtual maze task provides a more sensitive measure of trust towards specific targets than existing questionnaires or economic games. In a third experiment we demonstrate how the VR task could be adapted for traditional displays. We go on to use the maze task in Chapter 4 to test the effects of mimicry on trust.

2. Does being mimicked by a virtual character lead to rapport and trust?

Chapters 3 and 4 use virtual mimicry to strictly test the claim that mimicry leads to rapport and trust. In both studies we programmed virtual characters to mimic the head and torso movements of a participant with a specified time delay. In Chapter 3, we explored a range of outcomes reported in the previous literature and found that mimicry had a significantly positive effect on rapport. In Chapter 4, we carried out a more rigorous experiment in order to see if this effect could be replicated following a pre-registered design.

3. Are the effects of being mimicked modulated by the timing of mimicry?

Chapter 3 also exploits the advantages of virtual mimicry to test the role of timing. We showed that the timing of mimicry had a significant impact on conscious detection, but did not affect the strength of participants' positive responses to being mimicked.

4. Does mimicry lead to rapport and trust across group boundaries?

Chapter 4 investigates the effects of virtual mimicry within in-group and out-group pairs. The group membership of the mimicker did not significantly change levels of rapport or trust in response to being mimicked.

5. What are the parameters of mimicry in natural conversations?

After finding weak effects in the virtual mimicry studies, we depart from strict tests with virtual mimickers to investigate mimicry as it naturally occurs. Chapter 5 investigates the natural parameters of interpersonal coordination in face-to-face conversations between naïve participants. We used wavelet analysis to examine the dyadic coordination of head movements across a range of motion frequencies. Across two datasets, participants showed significantly less coordination than chance at head motion frequencies between 1.5 and 5 Hz, as well as greater than chance coordination at lower frequency ranges which would traditionally associated with mimicry timescales.

Chapter 2. The virtual maze: A behavioural tool for measuring trust

2.1 Abstract

Trusting another person may depend on our level of generalised trust in others, as well as perceptions of that specific person's trustworthiness. However, many studies measuring trust outcomes have not discussed generalised versus specific trust. To measure specific trust in others, we developed a novel behavioural task. Participants navigate a virtual maze and make a series of decisions about how to proceed. Before each decision, they may ask for advice from two virtual characters they have briefly interviewed earlier. We manipulated the virtual characters' trustworthiness during the interview phase and measured how often participants approached and followed advice from each character. We also measured trust through ratings and an investment game. Across three studies we found participants followed advice from a trustworthy character significantly more than an untrustworthy character, demonstrating the validity of the maze task. Behaviour in the virtual maze reflected specific trust rather than generalised trust, whereas the investment game picked up on generalised trust as well as specific trust. Our data suggests the virtual maze task may provide an alternative behavioural approach to measuring specific trust in future research, and we demonstrate how the task may be used in traditional laboratories.

2.2 Introduction

In this chapter, we aimed to develop a novel task for measuring trust towards specific people. In everyday life, we often have to weigh up how much we trust strangers or people we have only briefly met. Can I trust that passer-by to direct me to the station? If I lend the new intern my stapler, will I get it back? In some situations, such as criminal investigations, there are high stakes attached to the decision whether or not to trust someone. A wide body of literature suggests that nonverbal cues influence how trustworthy we perceive someone to be (Anderson, DePaulo, Ansfield, Tickle, & Green, 1999; DePaulo et al., 2003; Hosman & Wright, 1987). In the field of mimicry, it has been found that being mimicked can lead to persuasion and conformity (Leander et al., 2011; Tanner, 2008), although existing measures of trust have yielded mixed findings about the effect of being mimicked (Maddux et al., 2008; Verberne et al., 2013). In this chapter we consider how to measure the level of trust one person feels towards a specific stranger and develop a new behavioural task which may be used to test social factors such as mimicry that may influence levels of trust.

The measurement of trust is a vast and complicated topic, spanning domains in psychology, neuroscience, sociology, behavioural economics and organisation science (Ashraf, Bohnet, & Piankov, 2006; Bachmann & Zaheer, 2006; Blomqvist, 1997). Across these fields there are many different definitions of trust (Ashraf et al., 2006; Bachmann & Zaheer, 2006; Ben-Ner & Halldorsson, 2010), but very broadly speaking there are two main ways we can think about an individual's level of trust towards someone else. On one hand, we can treat their level of trust as a stable personal characteristic, a reflection of how much they trust others in general. This is often termed 'generalised trust' (Couch & Jones, 1997; Freitag & Traunmüller, 2009). On the other hand, we can treat their level of trust as a specific reaction to the other person, perhaps based on having a close relationship with them, or other social cues if the person is a stranger. The term 'interpersonal trust' is often used to refer to this kind of trust between people in a close relationship (Johnson-George & Swap, 1982; Rotter, 1967, 1971); we will use the term 'specific trust' to cover strangers as well.

In this chapter, we are interested in how to measure specific trust towards one particular person. We will begin by reviewing three major methods for trust measurement available to social psychologists and social neuroscientists, considering the suitability of each method. The first and oldest method is to administer questionnaires that ask people to self-report how much they trust others (Johnson-George & Swap, 1982; Rempel, Holmes, & Zanna, 1985; Rosenberg, 1957; Rotter, 1967). A more recent behavioural approach is to measure how much money an individual will entrust to another player in an economic game (Berg, Dickhaut, & McCabe, 1995; Glaeser, Laibson, Scheinkman, & Soutter, 2000; Tzieropoulos, 2013). Finally, an alternative behavioural approach is to measure trust in terms of willingness to ask for and endorse information from an informant (Clément, Koenig, & Harris, 2004; P. L. Harris & Corriveau, 2011). The first two approaches are the most widely used across psychology and other social sciences, but we argue the 'ask and endorse' approach is most suitable for measuring trust towards specific strangers. We then introduce our novel method for measuring trust using a virtual maze task, which builds on some of the advantages of ask-endorse paradigms.

2.2.1 Self-Report Questionnaires

For many decades, researchers have used self-report questionnaires to gauge levels of trust. Most of these have been designed to measure generalised trust (Couch & Jones, 1997; Freitag & Traunmüller, 2009; Glaeser, Laibson, Scheinkman, & Soutter, 1999) or interpersonal trust towards a close interaction partner such as a spouse or family member (Couch & Jones, 1997; Johnson-George & Swap, 1982; Rotter, 1967, 1971; Sorrentino, Holmes, Hanna, & Sharp, 1995). In comparison, there are very few validated scales that capture specific trust towards a stranger during an experiment. One exception is McCroskey & Teven's (1999) trustworthiness scale in which participants rate six dimensions of trustworthiness, in line with research highlighting the multidimensional nature of trustworthiness judgments and other social impressions (e.g. Ben-Ner & Halldorsson, 2010; Fiske, Cuddy, & Glick, 2007). Other researchers have often used just one or two items to measure the perceived trustworthiness of a target person (e.g. Maddux, Mullen, & Galinsky, 2008; Todorov, Pakrashi, & Oosterhof, 2009; Willis & Todorov, 2006). For example, participants might rate the target from 'not at all' to 'extremely' trustworthy (Willis & Todorov, 2006), or make a yes/no judgement about the target's trustworthiness and back this up with a confidence rating (Todorov et al., 2009). Approaches like these have the advantage that they are straightforward to administer and therefore may translate well to settings outside of the laboratory, particularly to clinical contexts in which behavioural measures may be impractical and time-consuming. However, one or two items may not fully capture the nature of trustworthiness (Ben-Ner & Halldorsson, 2010) and perceiving someone as trustworthy may not always equate to trusting them. It is also difficult to interpret and compare questionnaire results across studies, since the items used vary and are often specific to the experimental setting, e.g. ‘how much did you trust the other party during the negotiation?’ (Maddux et al., 2008). This can be an advantage for investigating context-specific trust towards different sources (e.g. political leaders vs. scientists on the topic of stem cell research; Liu & Priest, 2009) but is less suitable for social cognitive research into what makes people trust or not trust a specific stranger.

The lack of validated scales for measuring trust towards strangers may be due to concerns over the validity of self-report methods. Questionnaire items are open to interpretation, which may undermine the validity of responses. Trust questionnaires can be particularly susceptible to ambiguity due to the multiple meanings and interpretations of trust (Ben-Ner & Halldorsson, 2010; Lyon et al., 2012), especially across cultures (Hooghe, Reeskens, Stolle, & Trappers, 2009; Miller & Mitamura, 2003). Assuming the interpretation of items were unambiguous, participants still may not have accurate access to their internal feelings (Chan, 2009); even if they do, their self-reports may be biased by social norms or demand characteristics (McCambridge, de Bruin, & Witton, 2012). Finally, even if people report accurate internal feelings, these may be poor predictors of external behaviour (Armitage & Christian, 2003); for example, survey measures of generalised trust actually predict *trustworthy* behaviour better than trusting behaviour (Glaeser et al., 1999). For these reasons, implicit behavioural measures are often preferred over explicit self-report ratings.

2.2.2 The Investment Game

A major behavioural method for investigating trust emerged from behavioural economics. This method is aligned with the view that trust involves a ‘voluntary transfer of a good or favour to someone else, with future reciprocation expected but not guaranteed' (Gunnthorsdottir et al., 2002, p. 50). The method was developed by Berg, Dickhaut and McCabe (1995), who designed a simple investment game played between two people: an investor and a trustee. The investor was given $10 (different amounts have been used in subsequent studies; Johnson & Mislin, 2011) and had to decide how much of that $10 to send to the trustee, knowing that the amount they sent would be tripled before it was given to the trustee. Then the trustee had to decide how much of the tripled amount to return to the investor. The game measures trust behaviour in terms of the percentage of money the investor is willing to send to the trustee. Earlier trust games, such as the prisoners' dilemma, typically required an all-or-nothing decision to trust or distrust the other player, which did not provide such a sensitive measure (Schniter, Sheremeta, & Shields, 2013). Thus, the investment game has come to dominate the field (Johnson & Mislin, 2011).

It is unclear to what extent the investment game taps into generalised trust versus specific trust. Although the investment game was originally designed to answer questions about generalised trust (Berg et al., 1995; Glaeser et al., 2000; McEvily, Radzevick, & Weber, 2012), several studies have found that people's investments correlate poorly with generalised trust questionnaires (Ashraf et al., 2006; Glaeser et al., 1999) and relate instead to perceptions of the other player's trustworthiness (McEvily et al., 2012). In psychology and neuroscience, the investment game has been used to test factors that may affect specific trust towards known and unknown trustees. For example, studies have found that participants make significantly higher investments when the trustee is happy (Tortosa, Lupiáñez, & Ruz, 2013; Tortosa, Strizhko, Capizzi, & Ruz, 2013), belongs to a racial in-group (Stanley et al., 2012) or coordinates their nonverbal behaviour with the investor (Launay, Dean, & Bailes, 2013; Verberne et al., 2013, 2015). On the other hand, investment behaviour is also found to correlate with traits such as altruism (Ashraf et al., 2006; Ben-Ner & Halldorsson, 2010; Cox, 2004) and risk-seeking (Karlan, 2005; McEvily et al., 2012; Schechter, 2007), and to vary according to cultural norms (Johnson & Mislin, 2011; Willinger, Keser, Lohmann, & Usunier, 2003). This suggests that even if the investment game is not correlated with self-reported levels of generalised trust, it is sensitive to stable individual characteristics which may be proxies of generalised trust. Overall, the amount someone invests is likely to reflect a mixture of generalised trust and specific trust towards the other player, but it is unclear how levels of these are weighted in different people and different versions of the investment game.

The investment game also has some practical limitations. Firstly, it is hard to know how far people's investment behaviour in this abstract game may be used to infer how they would trust someone in the real world. In the investment game, participants have to make an explicit decision about the amount of money to send the trustee, but in real life people's trust decisions may be more implicit. Secondly, the rules of the investment game are somewhat complicated to explain, causing differences in how participants perceive and interpret the game (Macko, Malawski, & Tyszka, 2014). The task complexity means that the investment game may not be suitable for young children, or participant groups whose understanding is otherwise impaired. Furthermore, healthy adult participants may 'overthink' their response; for example, participants make more cautious investments when they have less time to make a decision (Tzieropoulos, Grave De Peralta, Bossaerts, & Gonzalez Andino, 2011) or know they will be paid randomly at the end of the study (Johnson & Mislin, 2011). Finally, the investment game is dyadic and over successive rounds players learn about each other's trust or trustworthiness (King-Casas et al., 2005). This means there is a one-shot opportunity on the first round to measure the investor's initial trust towards the trustee. It is unclear whether an averaged measure may be derived across multiple trials of the 'first round' where the investor does not find out the trustee's decision.

Economic trust games in virtual reality. In the context of investigating specific trust towards virtual characters, previous research has often construed trust in terms of cooperative behaviour (de Melo et al., 2011, 2013, 2012; de Melo, Zheng, & Gratch, 2009; Kulms, Kopp, & Krämer, 2014; Kulms, Mattar, & Kopp, 2015). In particular, several studies have investigated trust towards a virtual character in terms of willingness to cooperate rather than defect over the course of an iterated prisoners' dilemma game. For example, de Melo, Zheng & Gratch (2009) found that participants cooperated significantly more with a virtual character that expressed facial emotions after each round of the game (e.g. gratitude or reproach) compared to a virtual character with a neutral expression. In a more recent study (de Melo, Carnevale & Gratch, 2013) they found that participants cooperated more with an agent or avatar that made emotional expressions signalling cooperativeness (e.g. smiling after cooperation) compared to an agent or avatar that signalled competitiveness (e.g. smiling after exploiting the participant). Participants also rated cooperative agents or avatars as more trustworthy than competitive agent or avatars, suggesting that their behaviour in the prisoners' dilemma may be interpreted as trusting the other player (or not), in line with the use of this task in the economic literature. Very similar findings have been reported by this group (de Melo, Carnevale & Gratch, 2011; de Melo, Carnevale, Gratch & Read, 2012) and others (e.g. Kulms, Kopp & Krämer, 2014) when they recast the prisoners' dilemma game as an investment game, in which the choice to cooperate or defect was reframed as a choice to invest in a particular project, e.g. 'Project Green' vs. 'Project Blue'. Overall, this body of research suggests that the prisoners' dilemma (or variations thereof) can be a sensitive behavioural tool for measuring cooperation towards specific virtual characters presented sequentially in the same experiment.

On the other hand, Antos, de Melo, Gratch and Grosz (2011) found more mixed results using a version of the investment game (note that they refer to this as a variation of the public goods game, but their task actually follows the format of the investment game developed by Berg, Dickhaut & McCabe, 1995). In their study, participants completed a negotiation task with successive virtual characters that varied in their emotional expressions (five different conditions) and negotiation strategies (four different conditions) and accrued coins from successful negotiating. After every three negotiations, participants played a version of the investment game involving two steps: in the first step they had to choose which character from the preceding neogtiations they wanted to play with; in the second step they exchanged chose how many coins to invest with the chosen character following the typical investment game procedure. The results showed that the expression of the character and its strategy during the negotiation phase had significant interactive effects on which character the participant chose to play with. However, there were no significant effects on the amount of resources that participants trusted to the virtual character. This suggests that perhaps the binary selection part of the task provides a more sensitive measure than the transfer of resources. In the next section we elaborate on the potential advantages of tasks that use a binary selection to assess trust towards specific people.

In our own pilot experiments, we had little success using the investment game to test specific trust towards mimicking virtual characters. Across three pilot studies with a total of 78 participants we have consistently found that how much people invest with a particular virtual character does not correlate with their rating of that character's trustworthiness. Instead, we find that their investments towards two different virtual characters are highly correlated, suggesting that the investment game was picking up on a stable level of generalised trust. This motivated us to develop an alternative behavioural approach that would be more sensitive to levels of specific trust and less affected by generalised trust or other individual traits.

2.2.3 The Ask-Endorse Paradigm

An alternative behavioural approach to measuring trust comes from developmental psychology. In order to investigate the extent to which children will trust information from a teacher or informant, Koenig et al. (Koenig, Clément, & Harris, 2004; Koenig & Echols, 2003; Koenig & Harris, 2005, 2007; Pasquini, Corriveau, Koenig, & Harris, 2007) developed a paradigm in which children implicitly had to consider the trustworthiness of two puppets or video characters. Having seen one puppet or video character give accurate information, and the other give inaccurate information, each child had to make two choices. Firstly, they had to choose which one they would ask in order to learn something new. Each puppet or character would give conflicting testimony. Then the child would have to choose which testimony they believed. This procedure was repeated over multiple trials. Thus the paradigm provides two measures of trust: firstly, whether the child would *ask* for information, and secondly whether they would *endorse* that information. These may be called ‘selective’ trust measures since trust is inferred from which informant the child selects. Other selective trust research by Mills et al. (Johnston, Mills, & Landrum, 2015; Landrum, Mills, & Johnston, 2013; Mills, Legare, Bills, & Mejias, 2010) has also measured how often children would ask questions to different puppet informants to help them figure out a puzzle. Studies using selective trust paradigms have shown that children tend to trust informants who are nice, smart and honest (Landrum, Eaves Jr, & Shafto, 2015; Lane, Wellman, & Gelman, 2013), as well as people who are attractive (Bascandziev & Harris, 2014) or belong to their in-group (Kinzler, Corriveau, & Harris, 2011; VanderBorght & Jaswal, 2009). Several of the effects found with children have been replicated in adults using the same approach (Landrum et al., 2015; Lane et al., 2013).

There are several advantages to this approach for measuring trust towards a specific stranger. Firstly, it represents an ecologically valid scenario, i.e. based on limited experience of this person, will you ask them for advice and trust what they say? We might consider the same questions when asking a passer-by for directions, or conducting an investigative interview. Therefore, the ask-endorse scenario might be more representative of everyday trust decisions than the investment game scenario. Secondly, the decision to ask or endorse taps into implicit trust behaviour, which may provide a ‘purer’ estimate of trust levels than explicit measures. Thirdly, by framing each decision in terms of selecting one informant versus the other, participants have to use their perceptions of each informant’s trustworthiness and therefore the task should be more sensitive to specific trust than generalised trust.

2.2.4 The Virtual Maze

We aimed to design a novel behavioural task for measuring trust towards specific strangers, and decided to implement the task in virtual reality. As virtual reality technology becomes increasingly sophisticated and more widely available, it is becoming an increasingly popular tool in social psychology and social neuroscience (Chapter 1; Blascovich et al., 2002; McCall & Blascovich, 2009). This is because virtual characters and virtual spaces can be manipulated in systematic ways that would be hard to achieve with confederates in a physical laboratory (Fox et al., 2009). For example, it is easy to program a virtual character that subtly blinks or nods at a certain rate (Bailenson & Yee, 2005; Gratch et al., 2006), or a virtual space which is much larger than the average laboratory. Many studies have shown that participants react to virtual scenarios as they would in real life (Durlach & Slater, 2000; Fox et al., 2009; Garau et al., 2005), so virtual reality also offers an opportunity to closely replicate everyday situations under controlled conditions. As well as affording high ecological validity, there are also novel opportunities to measure implicit behavioural responses such as where participants direct their gaze in the 3D space, or how closely they will approach a virtual character (Khooshabeh et al., 2011; McCall et al., 2009).

Exploiting some of the strengths of virtual reality, we designed a task where participants navigate through a virtual maze and may choose to trust virtual characters about which way to proceed. Participants find themselves in a virtual maze made up of a series of identical rooms connected by corridors. Each time they enter a new room, they face two doors and must make a decision about which door to proceed through. To help them decide, they may approach two virtual characters for advice, although they do not have to. When approached, each virtual character will indicate which door they think is the one to take. The participant keeps making decisions until they are told that they have found the way out of the maze. In fact, there are no right or wrong choices about which way to go. Instead, we are able to randomly generate endless rooms and corridors until the participant has gone through a specified number of rooms (trials) and we tell them they have found the way out. At the end, we can measure how often the participant approached each character for advice, and how often they followed advice they received from each character. The virtual maze therefore follows a similar approach to the ask-endorse paradigm, although the virtual maze task is more implicit in that participants are not prompted on every trial which character they want to ask for advice, but are instead left to make an implicit choice to approach neither, one or both characters.

In this chapter, we present three studies in which we piloted the virtual maze task and explored how it compares to other trust measures. Firstly, we aimed to test whether people's decisions in the virtual maze are sensitive to differences in trustworthiness between two virtual characters. Therefore, in the first two studies participants got to know two different virtual characters through a short interview where the participant asked each character some prepared questions. During the interview, we manipulated the trustworthiness of each character through their verbal answers and nonverbal and vocal behaviour. The manipulation was designed to achieve a large effect size, since we did not know the sensitivity of the virtual maze task. We predicted that participants would decide to approach and follow the advice of a trustworthy character significantly more often than an untrustworthy character in the virtual maze. In order to compare our task with major alternative methods, in Studies 1 and 2 we also included ratings of each character's trustworthiness and an investment game. We further aimed to explore the extent to which each measure showed a correlation between the two virtual characters. If trust in one character is correlated with trust in the other character, that would indicate the level of trust towards each one is being driven by stable individual differences among participants, i.e. generalised trust. On the other hand, if a measure shows no correlation between two characters that suggests it is sensitive to specific trust rather than generalised trust. In the third study, we implemented a purer manipulation of trustworthiness by using the investment game as a manipulation instead of a dependent measure, following two studies by Franzen et al. (Franzen et al., 2011; Lis et al., 2013). In Study 3, we also programmed a low-tech version of the virtual maze task on a standard computer screen to demonstrate how the task could be used in standard laboratories without VR equipment or software.

2.3 Study 1

2.3.1 Methods

Participants. Twenty participants (13 female) were recruited through email advertisements to the Institute of Cognitive Neuroscience (ICN) departmental participant pool. A power calculation using G*Power (version 3.1.7, 2013) showed that 20 participants would be sufficient to detect a large effect size (d_z = 1.1) with power of 0.98 or a medium effect size (d_z = 0.71) with power of 0.91. All participants gave written informed consent and received payment of £7, plus a bonus of up to £3 depending on how they played the investment game. The study was granted ethical approval by the MoD Research Ethics Committee (Reference Number 564/MODREC/14) and the ICN Ethics Chair (Project ID Number ICN-AH-PWB-3-3-14a).

Virtual reality system. Participants sat at a desk in front of a 90 x 160cm projector screen. We used Vizard virtual reality software (Worldviz, 2014) to display a virtual environment on the screen (Figure 2-1A). During the interview phase of the experiment, the virtual environment looked like an extension of the physical desk and walls of the laboratory. We programmed life-sized virtual characters to appear seated on the other side of the desk, facing the participant. We programmed a virtual barrier to occlude the virtual character at the end of the interview phase. Instructions and stimuli for the following tasks were displayed on the virtual barrier. The virtual maze task was also completed on the projector screen, but the maze environment was not designed to look like an extension of the laboratory. Instead, the participant saw virtual corridors and rooms similar to playing a video game.

Virtual characters. We prepared two male virtual characters for the experiment, named Mike and Ryan. The characters' appearances (Figure 2-1B) were selected from a collection of characters ('Complete Characters HD') supplied by Rocketbox Libraries for Vizard. We scripted everything the characters said during the experiment, which we pre-recorded from two male volunteers with native British accents. The pre-recorded speech was triggered by the experimenter or the computer program. The characters were programmed to move their jaw according to the amplitude of the pre-recorded speech, so that it looked like they were speaking. Audio speakers were hidden behind the projector screen so that the sound of the character's voice came from their virtual location.

Trustworthiness manipulation. We manipulated the trustworthiness of the two virtual characters during the initial interview phase of the experiment. Mike was designed to seem trustworthy, whereas Ryan was designed to seem untrustworthy. We achieved the manipulation through verbal, non-verbal and vocal signals.

Verbal signals. Mike and Ryan gave different scripted responses to three interview questions: (1) *What is your occupation?* (2) *What did you do last weekend?* (3) *What are your plans for the summer?* Mike made statements demonstrating reliability, e.g. *'I promised to raise £800 in sponsorship and I managed to smash the target'*, whereas Ryan indicated irresponsibility, e.g. *'things didn't go so well at my last job and I basically ended up getting fired after missing too many deadlines'*. For full scripts see Appendix.

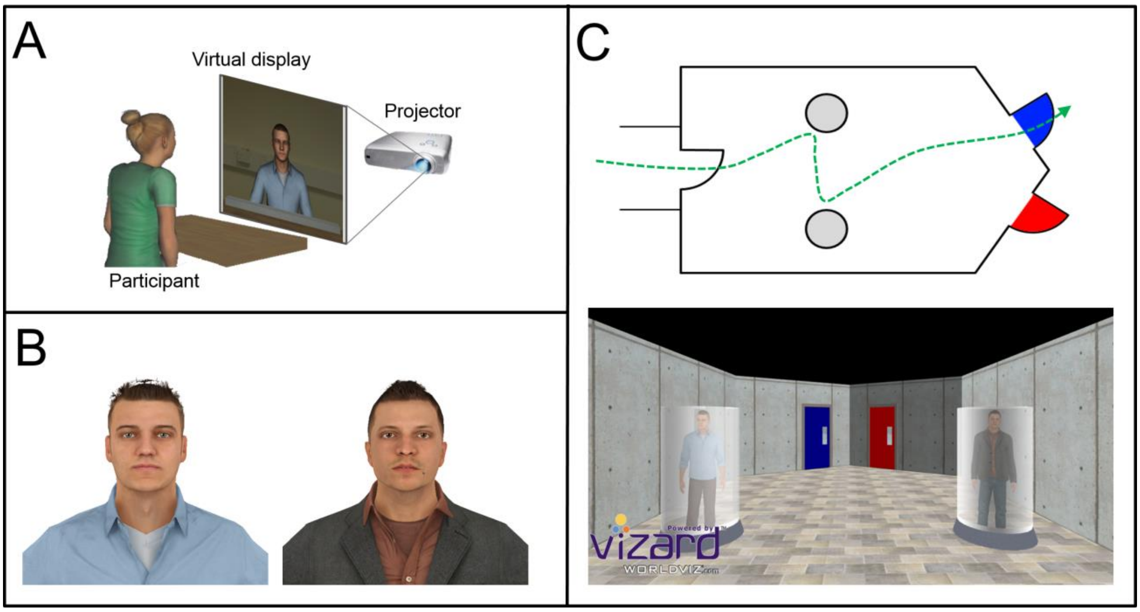

Figure 2-1. Overview of the virtual display (Study 1). Participants interacted with virtual characters displayed on a projector screen (panel A). Panel B shows the appearance of Mike (left) and Ryan (right). Panel C shows a plan and screenshot of each room in the virtual maze. The dashed green line illustrates a possible path through the room.

Non-verbal signals. We programmed Mike and Ryan to make different amounts of eye contact, which people often use as a nonverbal signal to trustworthiness (even if it is not a reliable signal; DePaulo et al., 2003). Mike averted his gaze at random intervals of 6-9s, whereas Ryan averted gaze at random intervals of 3-6s and consequently made less eye contact. Both characters looked away at a random location for 0.75-1.75s before returning to make eye contact.

Facial appearance is also used as a nonverbal cue to trustworthiness (Todorov, 2008; Todorov & Duchaine, 2008). We obtained pilot ratings of trustworthiness on a 1 to 6 scale for six different virtual characters by 52 participants in an online questionnaire. The appearance rated the most trustworthy (M = 3.89, SD = 1.15) was assigned to Mike and the appearance rated as least trustworthy (M = 2.79, SD = 1.16) was assigned to Ryan (c.f. Ewing, Caulfield, Read, & Rhodes, 2014). There was a significant difference in trustworthiness ratings between the two characters ($t(51) = 5.37$, $p < .001$).

All other nonverbal behaviour (e.g. posture, head tilting) was controlled using the same idle animation for both characters.

Vocal signals. Vocal hesitations and disfluencies are sometimes seen as a signal of untrustworthiness (Anderson et al., 1999; Hosman & Wright, 1987). Therefore, for the interview phase (but not the virtual maze task), the volunteer voicing Mike was instructed to speak clearly without many hesitations, whereas the volunteer voicing Ryan was instructed to mumble and make regular hesitations (e.g. *'umm'*).

Virtual maze task.

Virtual maze environment. The virtual maze was generated from a series of identical rooms (Figure 2-1C). The rooms were connected by twisting sections of corridor designed to enhance the illusion of being in a maze. The participant entered each room through a brown door; at the far end of the room, there was a red door and a blue door. In each room there were also two semi-transparent 'hologram chambers', where the virtual characters appeared as holograms from outside the maze. Mike always appeared on the left and Ryan always appeared on the right, but participants did not show any preference for approaching the left versus right hologram chamber (see Appendix). Whenever the participant got close to a hologram chamber, there was a sound effect and the chamber became more transparent. At the same time, the character inside the chamber would spin to face the participant and deliver some verbal advice. Participants were able to navigate the virtual environment by using a keypad to turn the viewpoint and move forwards or backwards.

Trial procedure. Going through one room corresponded to completing one trial. Each participant completed twelve trials in total. In each trial, they had to make a choice about whether to proceed through the blue door or the red door. The participant was able to approach neither, one or both characters to receive advice about which door to choose, although this was not explicitly instructed. If approached, the characters randomly delivered uncertain advice about which door the participant should choose (e.g. *'It's blue this time, I think'*). There was no 'correct' door on each trial. Instead, the maze was completed after twelve trials in which the participant approached at least one character. If a participant asked neither character, that trial was recorded but did not count towards the requisite twelve trials. This ensured we had twelve trials in which the participant received some advice about which way to go, thus providing data about how much they trusted that advice.

Character advice. Mike and Ryan were programmed to advise the red door in half the trials and the blue door in the other half. They were also programmed to advise the same door as each other in half the trials and different doors in the other half. In order to generate the verbal advice stimuli, we pre-recorded twelve scripted phrases. Then we paired the phrases in order to create twelve combinations of advice stimuli (Table 2-1). The order of the stimulus combinations was randomised for each participant. Note that the participant would only receive an advice stimulus if they approached a character for advice. Therefore, some participants may not have received both parts of every stimulus combination.

Dependent measures. We measured three dependent variables on each trial of the maze task. Firstly, we recorded whether each character was approached for advice. Secondly, we recorded which character the participant approached first (if the participant only approached one character, then that character was treated as being approached first). Thirdly, we recorded whether the participant followed the advice of each character. We averaged each of these variables over the total number of trials completed, giving us (1) the proportion of trials on which each character was approached for advice; (2) the proportion of trials on which each character was the first (or only) character approached for advice; and (3) the proportion of trials on which the participant followed each character's advice.

Table 2-1. Advice stimuli.

Trial	Advice from Mike	Advice from Ryan
1	I think you should try the blue door.	I think you should go through the red door.
2	I think you go through the blue door this time.	I think it's the red door this time.
3	It's the blue door, I think.	It's red this time, I think.
4	I think you should try the red door.	I think you should go through the blue door.
5	I think you go through the red door this time.	I think it's the blue door this time.
6	It's the red door, I think.	It's blue this time, I think.
7	It's blue this time, I think.	It's the blue door, I think.
8	I think it's the blue door this time.	I think you go through the blue door this time.
9	I think you should go through the blue door.	I think you should try the blue door.
10	It's red this time, I think.	It's the red door, I think.
11	I think it's the red door this time.	I think you go through the red door this time.
12	I think you should go through the red door.	I think you should try the red door.

Procedure. Kathryn Taylor carried out the data collection, supervised by Joanna Hale.

Interview. Participants were instructed that they were going to interview an individual virtual character called Mike (or Ryan) in order to get to know him. The participant was given a sheet with three prepared questions, and instructed to ask each question one at a time during the interview. At the start of the interview, Mike introduced himself and prompted the participant to begin asking questions. When the participant asked a question, the researcher triggered a pre-recorded scripted answer. The interview lasted around 5 minutes. At the end of the interview a virtual barrier appeared to occlude Mike and present the next tasks.

Character ratings. Participants rated their agreement with 10 statements about Mike by clicking on a continuous scale from 0 (strongly disagree) to 1 (strongly agree). Each statement began *'I think Mike is...'* followed by five items measuring rapport (*likable/engaging/kind/unfriendly/unpleasant*) and five items measuring trustworthiness (*trustworthy/honest/responsible/unreliable/insincere*). The statements were presented in a randomised order. We reversed the scores for negatively valenced items and averaged the responses to provide one score for rapport (α =.86) and one score for trustworthiness (α = .76).

Investment game. Participants completed five trials of the Investment Game with Mike, based on Berg et al.'s (1995) paradigm. At the start of each trial the participant had £1 to play with (this did not accumulate over trials). They could invest any proportion of the £1 with Mike. Mike would always triple their investment. Then he would choose a proportion of the tripled amount and return it to the participant, but his decision was not revealed. The participant would end the trial with any money they chose not to invest, plus any money that Mike chose to return them. Unknown to the participant, we programmed the task so that 50% of the tripled money was always returned (i.e. a maximum of £1.50 from Mike or Ryan). We told participants at the start of the game that one trial would be selected at random and they would find out the outcome of that trial, which would be paid as a cash bonus. To measure how much the participant trusted Mike, we measured the proportion of £1 that the participant chose to invest as an average across the five trials.

Once the participant had completed the interview, character ratings and investment game with Mike, then they completed the same three steps with Ryan. The order in which participants met Mike and Ryan was counterbalanced.

Virtual maze. After interacting separately with Mike and Ryan, the participant completed the virtual maze task. First the participant practiced using a keypad to navigate around a virtual space and approach hologram chambers in order to receive a greeting message from Mike or Ryan. After the practice, the participant was instructed that their task was to find the way out of the virtual maze as quickly as possible. They were also told ‘There are some rooms in the maze where you will have to make a choice about which way to go. To help you decide, Mike and Ryan will be able to give you remote advice from outside the maze. They will appear as holograms in the room and you can go up to them to get advice’. Then participants began the task. After completing twelve trials, they were told they had successfully completed the maze.

Finally, participants completed a questionnaire unrelated to the current study, and then received payment. The experiment took approximately 40 minutes.

2.3.2 Contributions

Joanna Hale wrote the interview scripts and programmed the virtual characters’ behaviour. Antonia Hamilton and Joanna Hale devised the maze task, which was programmed by Joanna Hale. Kathryn Taylor carried out participant recruitment and data collection, under the supervision of Joanna Hale. Joanna Hale carried out the analyses that follow.

2.3.3 Results and Discussion

Missing data. One participant completed only five trials of the virtual maze due to motion sickness, but their data was included in the analyses. Three other participants briefly paused during the virtual maze task due to feeling motion sick, but this was not enough to cause concern that the task was too unpleasant.

Trust towards Mike and Ryan. We carried out paired-samples *t*-tests to determine whether there was a significant difference between Mike and Ryan on each of our six dependent measures: rapport rating, trustworthiness rating, percentage investment, approaching for maze advice, approaching first, and following maze advice. We applied a Bonferroni correction for multiple comparisons. Mike scored higher than Ryan on every measure, indicating participants liked and trusted Mike more than Ryan (Table 2-2). In the virtual maze participants approached Mike significantly more often than Ryan, although they did not approach Mike first significantly more often than they approached Ryan first. Furthermore, participants followed advice from Mike significantly more than Ryan. These results show that the maze task was sensitive to our manipulation, although it is possible that the effects were due to differences in the likeability of each character, rather than their trustworthiness.

Table 2-2. Differences between Mike and Ryan (Study 1).

Measure	Mike *M* (*SD*)	Ryan *M* (*SD*)	Difference		
			t(19)	*p*	*d*
Rapport rating	.88 (.12)	.39 (.24)	8.28	< .001	2.58
Trustworthiness rating	.85 (.11)	.31 (.22)	8.59	< .001	3.10
Investment	.76 (.24)	.46 (.36)	6.29	< .001	0.98
Approach for advice	.90 (.18)	.62 (.35)	3.27	.004	1.01
Approach first	.62 (.27)	.38 (.27)	1.94	.07	0.89
Follow advice	.79 (.20)	.42 (.25)	4.32	<.001	1.63

Correlations between Mike and Ryan. For each measure, we examined the correlations between scores for Mike and Ryan (Table 2-3). Investment was the only measure which showed a significant correlation between the two characters, which was strongly positive. This means that participants were fairly consistent in their investments towards the two characters, suggesting investment behaviour might have been driven by generalised rather than their perception of the specific characters' trustworthiness.

However, we cannot tell whether this evidence for generalised trust may be explained by variability in traits such as risk aversion and altruism (McEvily et al., 2012), or whether different participants simply interpreted the task instructions differently (Macko et al., 2014).

Summary. We found that the maze task was sensitive to our trustworthiness manipulation, showing a significant difference in how often people would approach and follow the advice of the trustworthy virtual character (Mike) versus the untrustworthy character (Ryan). Our manipulation led to large effect sizes on the measures we took from the virtual maze, as well as our other dependent measures. The investment game was the only measure for which people's responses towards Mike and Ryan were significantly correlated. This suggests that people's investments may reflect a generalised level of trust, or other stable individual characteristics.

Table 2-3. Correlations between Mike and Ryan (Study 1).

Measure	Correlation	
	***r*(18)**	***p***
Rapport rating	.02	.93
Trustworthiness rating	-.38	.10
Investment	.82	< .001
Approach for advice	.10	.67
Approach first	-1.0*	0
Follow advice	-.23	.20

*Note that participants could only approach one character first on each trial and therefore the 'approach first' measure is perfectly negatively correlated across characters.

2.4 Study 2

In Study 2, we aimed to replicate Study 1 using a more controlled manipulation of trustworthiness, where each character would be matched in likeability. We also aimed to adapt the maze task for use with a head-mounted display (HMD), and included an extra questionnaire to measure how real the virtual experience seemed to participants and record any feelings of motion sickness associated with the HMD.

2.4.1 Methods

Participants. Twenty-four participants (17 female, M_{age} = 25.9, SD_{age} = 10) were recruited through email advertisements to the ICN departmental participant pool. The sample size was based on Study 1, which found large effects with a sample of 20. All participants gave written informed consent and received payment of £7 per hour, plus a bonus of up to £3 depending on how they played the investment game. The study was granted ethical approval by the MoD Research Ethics Committee (Reference Number 564/MODREC/14) and the ICN Ethics Chair (Project ID Number ICN-AH-PWB-3-3-14a).

Virtual reality system. We used Vizard virtual reality software (Worldviz, 2014) to display virtual environments in an Oculus Rift DK2 head-mounted display (HMD). This device allows people to look around a virtual 3D space as if they are really there. Participants wore the HMD while seated at a physical desk in our lab. During the interview phase of the experiment, participants saw a virtual desk in place of the physical desk, inside a virtual room that looked like a typical psychology laboratory. We programmed a virtual screen to occlude the virtual character at the end of the interview phase and display the next tasks to the participant. The virtual maze task was also displayed via the HMD. Participants wore stereo headphones throughout the experiment to hear the characters speaking and other sound effects. They were provided with a joystick to make responses during the tasks and navigate through the virtual maze.

Trustworthiness manipulation. The Mike and Ryan characters from Study 1 were slightly modified for Study 2. In this study, Mike was designed to seem reliable and Ryan was designed to seem unreliable but we aimed to make both characters equally likeable. Therefore, we scripted new questions and responses for the characters to deliver in the interview phase. We altered verbal, non-verbal and vocal signals of trustworthiness from each character as follows:

Verbal. Mike and Ryan delivered new scripted responses to seven interview questions. Mike made statements demonstrating reliability, e.g. *'I ended up graduating with the highest grade'*, whereas Ryan indicated unreliability, e.g. *'I only had a few lectures a week and I used to miss them all the time'*. Both characters showed their likeability, e.g. *'I met so many people that I'm still really close with'* (Mike); '*I also get along with everyone'* (Ryan). For full scripts see Appendix. To validate the new scripts, a separate pilot sample provided online ratings of rapport and trust towards the speakers. We presented the scripts from Study 1 and Study 2 as if they were transcripts from real interviews (in the scripts from Study 1 we changed the characters' names to David and Ben). Pilot participants rated significantly more trust (t(14) = 8.14, p < .001, d = 3.73) and rapport (t(14) = 8.46, p < .001, d = 2.16) towards Mike than Ryan in the Study 1 scripts (Table 2-4). In the new scripts for Study 2 (Table 4), pilot participants rated significantly greater trust towards Mike (t(14) = 3.56, p = .003, d = 1.48) but there was no significant difference in rapport towards Mike and Ryan (t(14) = .67, p = .514, d = 0.24), supporting the validity of our manipulation. However, it should be noted these pilot ratings were based on a small sample size.

Table 2-4. Descriptive statistics for each script (Study 2).

Study	Character name (as in online questionnaire)	Character trustworthiness	Character likeability	Rapport *M* (*SD*)	Trust *M* (*SD*)
1	Mike (David)	High	High	4.37 (.85)	4.83 (.77)
	Ryan (Ben)	Low	Low	2.23 (1.11)	1.67 (.92)
2	Mike (Mike)	High	High	4.93 (1.04)	4.71 (1.03)
	Ryan (Ryan)	Low	High	4.70 (.91)	3.29 (.88)

Non-verbal. As in Study 1, Mike and Ryan were programmed to make different amounts of eye contact. In this study, we also manipulated the characters' promptness as a signal of reliability. At the start of the interview phase, the participant was asked to wait while the current virtual character got ready. Mike was ready after a short delay of

4 seconds, saying '*Ok, yep, I'm ready to start!',* whereas Ryan took 14 seconds to get ready before saying 'S*orry I'm late. Yeah, OK, I'm ready now'.* All other nonverbal behaviour was controlled using the same idle animation for both characters.

Vocal. As in Study 1, the volunteer voicing Mike was instructed to speak without many hesitations, whereas the volunteer voicing Ryan was instructed to make more hesitations. However, both volunteers spoke clearly and engagingly so as to avoid differences in friendliness or likeability.

Virtual maze task. We slightly adapted the virtual maze task from Study 1 for use with the HMD. Instead of using a keypad to navigate the virtual maze, participants used a joystick. We triggered extra sound effects when the participant went through doors in the virtual maze, in order to increase realism. We also rendered a plainer texture on the virtual walls and simplified the sections of corridor joining each room in order to reduce motion sickness associated with navigating narrow and twisty sections of corridor. All other aspects of the virtual maze task remained the same.

VR Questionnaire. We included a short questionnaire about the participant's experience in immersive virtual reality. Participants indicated their agreement with statements on scale from 1 (strongly disagree) to 7 (strongly agree). Four statements assessed how much the virtual world was real ('presence'), e.g. *'During the experience, the interview felt like the real world for me'.* Four statements assessed how much the virtual characters seemed real ('co-presence'), e.g. *'My feelings and emotions in relation to Mike/Ryan were as if they were real'*. Items were based on presence and co-presence items used in other virtual reality studies (Friedman et al., 2014; Pan et al., 2012). The questionnaire also asked participants to indicate whether they felt any motion sickness, queasiness, headache or eye strain whilst wearing the HMD.

Procedure. Madeleine Payne carried out the data collection, supervised by J.H.

Each participant was seated at a desk and given verbal instructions from the experimenter. Then they were fitted with the HMD and saw themselves in a virtual laboratory. They could see a laptop on a desk in front of them, but a large screen displaying instructions blocked the participant's view of the room on the other side of the desk (Figure 2-2A). The participant was given a few minutes to become accustomed to the virtual environment. They also practiced using the joystick to trigger instructions and log responses. When the participant had completed the practice, they were instructed they were going to interview a character called Mike and were asked to press a button on the joystick when they were ready to start the interview. Then the participant was instructed to wait while Mike got ready. After a specified delay, the participant heard a door opening and a chair moving somewhere on the other side of the large screen, and then Mike said he was ready. At that point, the screen moved up to the ceiling so that the participant could see Mike (Figure 2-2B). Interview questions were displayed to the participant one at a time on the laptop. When the participant asked each question, the researcher triggered a pre-recorded scripted answer from Mike. The interview lasted around five minutes. At the end of the interview, the large screen was lowered again to occlude Mike and present the next tasks.

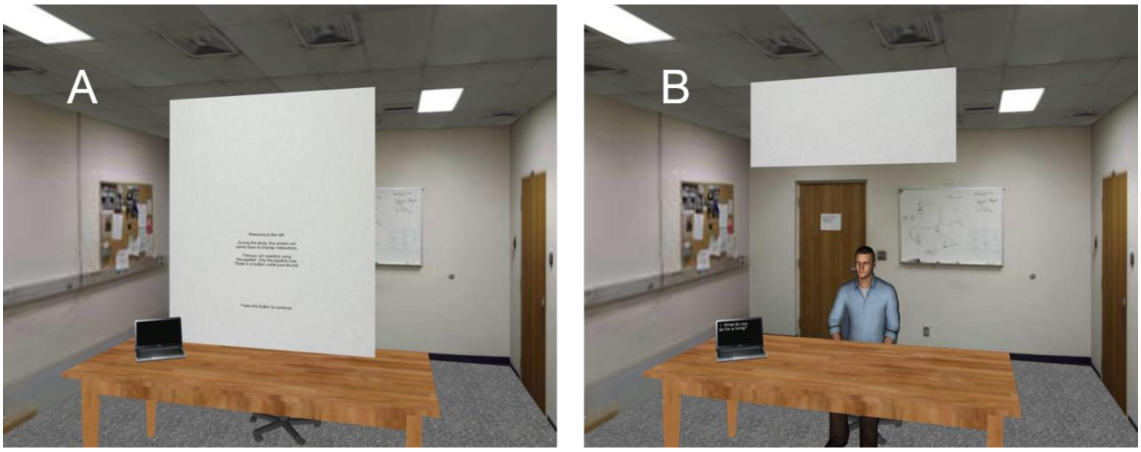

Figure 2-2. Virtual laboratory space (Study 2). The room is shown from the participant's side of the desk. Instructions and stimuli were displayed on large screen which occluded the virtual character (Panel A). During the interview phase, the participant saw the virtual character sitting opposite them and interview questions were displayed on the laptop screen (Panel B).

Character ratings. Immediately following the interview, participants rated their feelings of rapport and trust towards Mike, as in Study 1.

Investment Game. Next, participants completed the investment game as in Study 1. Data from Study 1 indicated that participants made highly consistent choices over five trials of the investment game. Therefore in Study 2 we decided to reduce the number of trials to one, consistent with the traditional paradigm (Berg et al., 1995). Reducing the number of trials also reduced the time participants spent wearing the HMD.

Once participants had completed these steps with Mike, they took a break from wearing the HMD and completed the VR questionnaire on paper. Then the same procedure was repeated with Ryan. The order of characters was counterbalanced across participants.

Virtual Maze. Then the participant completed the virtual maze task in the HMD. First, they practiced using the joystick to navigate around the virtual space and approach Mike and Ryan. They were instructed their task was to find the way out of the virtual maze through as few rooms as possible. They were also instructed ‘in each room in the maze, you must choose which way to go. To help you decide, Mike and Ryan can give you advice.’ Then participants began the task. After completing twelve trials, they were told they had successfully completed the maze.

Finally, participants received payment. The experiment took approximately 50 minutes.

2.4.2 Contributions

Joanna Hale programmed the virtual characters’ behaviour and the virtual maze task. Madeleine Payne wrote the interview scripts, recruited participants and carried out data collection, under the supervision of Joanna Hale. Joanna Hale carried out the analyses that follow.

2.4.3 Results and Discussion

Missing data. Three participants terminated the virtual maze due to discomfort before completing 12 trials. One participant reported making a mistake on two ratings;

these responses were recorded as missing. No participants were excluded from the analyses.

VR Questionnaire. The median rating of how real the virtual environment seemed to participants was 4.75 out of 7 (*M* = 4.58, *SD* = 0.73, *Range* = [2.75, 6]), and the median rating of how real the virtual characters seemed was 5 out of 7 (*M* = 4.72, *SD* = 1.25, *Range* = [1.5, 6.5]) which is very similar to levels found in other immersive VR experiments using a comparable rating scale (e.g. Friedman et al., 2014; Pan et al., 2012). One participant reported eye strain due to not wearing their glasses, but no other symptoms of motion sickness, queasiness, headache or eye strain were reported, despite three participants terminating the maze task early.

Trust towards Mike and Ryan. We carried out paired-samples *t*-tests to determine whether there was a significant difference between Mike and Ryan on each of our measures. We applied a Bonferroni correction for multiple comparisons. Replicating the results from Study 1, Mike scored significantly higher than Ryan on every measure (Table 2-5), indicating that participants liked and trusted Mike more than Ryan. In the virtual maze, participants approached Mike significantly more than Ryan overall, and Mike was also approached first significantly more than Ryan. Participants also followed advice from Mike significantly more than advice from Ryan. Therefore in Study 2 we replicated the results from Study 1 using a different trustworthiness manipulation and a more immersive virtual reality system.

The significant difference in rapport was inconsistent with our pilot ratings of Mike and Ryan's scripted responses, raising the possibility that our results were affected by other aspects of the experiment. In particular, the extra nonverbal cues which were present in the experiment (such as eye contact) could have affected feelings of rapport towards each character. Other factors such as the immersive environment could also have affected the experimental results. While we could not test all of these factors, we were able to test whether the order of interviewing Mike and Ryan (which was counterbalanced) affected our results. Because our experiment had a within-participants design, it is possible that participants' impressions of each character would differ depending on whether they already had a 'baseline' impression about the other character, or no baseline information to go on. Therefore we carried out ANOVAs to test the interaction between character (within-subjects) and interview order (between-subjects) on each of our dependent measures.

Table 2-5. Differences between Mike and Ryan (Study 2).

Measure	Mike *M* (*SD*)	Ryan *M* (*SD*)	Difference		
			***t*(23)**	***p***	***d***
Rapport rating	.82 (.12)	.60 (.16)	5.26	<.001	1.56
Trust rating	.74 (.17)	.42 (.13)	6.45	<.001	2.11
Investment	.77 (.25)	.48 (.36)	4.93	<.001	0.94
Approach for advice	.98 (.05)	.76 (.31)	3.67	.004	0.99
Approach first	.64 (.29)	.36 (.29)	2.38	.03	0.97
Follow advice	.87 (.15)	.47 (.23)	5.60	<.001	2.06

There was a significant interaction between character and interview order on rapport ratings ($F(1, 22) = 4.55$, p = .04, $n_p^2 = .17$) and investments ($F(1,22) = 5.47$, p = .03, $\eta_p^2 = .20$). To further decompose these effects, we conducted post-hoc pairwise comparisons with Bonferroni correction. For rapport, Mike was rated significantly higher than Ryan by both groups, but the difference was greater in the group who interviewed Mike first ($M_{Mike} = .83$, $M_{Ryan} = .53$, $t(10) = 4.76$, $p = .004$) compared to the group who interviewed Ryan first ($M_{Mike} = .81$ $M_{Ryan} = .67$, $t(12) = 3.11$, $p = .037$). For investments, only the group who interviewed Mike first invested significantly more money with Mike than Ryan ($M_{Mike} = .67$ $M_{Ryan} = .25$, $t(10) = 5.62$, $p < .001$); the same effect was not significant for participants who interviewed Ryan first ($M_{Mike} = .85$ $M_{Ryan} = 0.68$ $t(12) = 2.30$, $p = .16$). A graph of these results highlights the presence of an order effect (Figure 2-3). Participants invested almost identical amounts of money in the first character they interviewed, but when we look at the second character interviewed there is a large difference between Mike and Ryan. This suggests that participants' investment in the first character simply reflected their level of generalised trust, but once they had a baseline impression about the first character their investment with the second character was more informed by perceptions of trustworthiness, i.e. specific trust.

Correlations between Mike and Ryan. For each measure, we examined the correlations between scores for Mike and Ryan (Table 2-6). Replicating our results from Study 1, investment was the only measure which showed a significant positive correlation between the two characters. We interpret this as evidence that the investment game is sensitive to generalised trust or other stable individual characteristics, since people showed strong consistency in their investments towards different virtual characters. We also found a significant negative correlation between following Mike's advice and following Ryan's advice in the maze. This reflects the fact that on half of the trials in the virtual maze task, Mike and Ryan gave opposing advice.

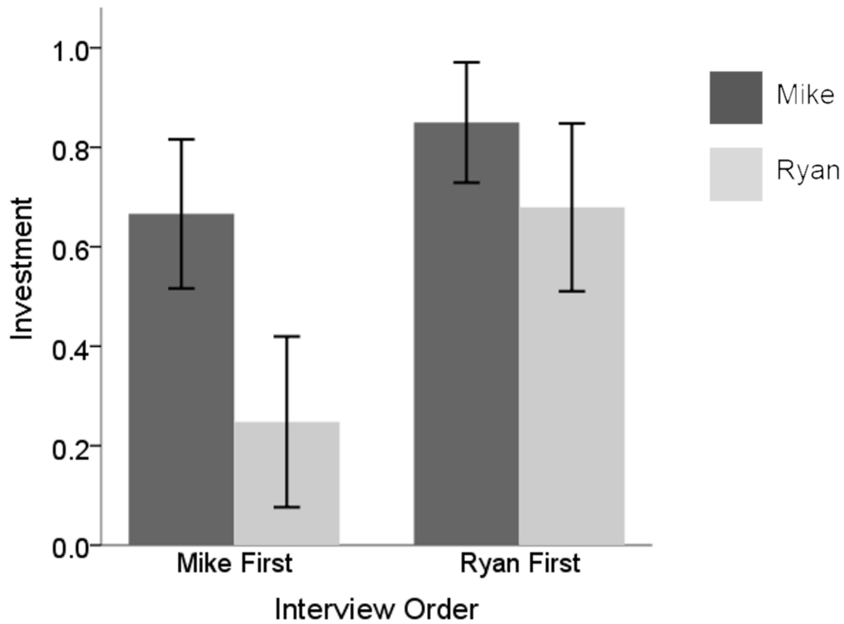

Figure 2-3. Investments towards Mike and Ryan by interview order (Study 2).

Table 2-6. Correlations between Mike and Ryan (Study 2).

Measure	Correlation	
	r(22)	*p*
Rapport rating	-.05	.80
Trust rating	-.23	.28
Investment	.63	.001
Approach for advice	.18	.40
Approach first	-1.0*	0
Follow advice	-.64	.001

*Note that participants could only approach one character first on each trial and therefore the ‘approach first’ measure is perfectly negatively correlated across characters.

Summary. We found that the virtual maze task was sensitive to our more controlled manipulation of trustworthiness, replicating the results from Study 1. Participants were significantly more likely to approach and follow the advice of the trustworthy virtual character (Mike) versus the untrustworthy character (Ryan). We also found a significant effect of our manipulation on the other measures of trust as well as on rapport ratings, despite the pilot ratings of our new scripts suggesting that Mike and Ryan were similarly likeable. Therefore, we explored the possibility of an order effect in our data and found that the order in which the virtual characters were interviewed affected how much participants invested in Mike versus Ryan. Those who interacted with Mike first invested much less money with Ryan compared to Mike, but those who met Ryan first invested similar amounts with Ryan and Mike. In fact, participants invested very similar amounts in whichever character they met first, but when investing with the second character they showed greater trust towards Mike than Ryan. This suggests that people’s initial investments may have been driven by generalised trust, but their investment with the second character may have been more informed by perceptions of trustworthiness, i.e. specific trust. This interpretation is consistent with our finding that investment was the only measure showing a significant positive correlation between the two characters, which replicated Study 1 and indicates that the investment game at least partly reflects stable individual differences.

2.5 Study 3

In the first two studies we demonstrated that the virtual maze task is sensitive to differences between two characters and appears to tap into specific trust rather than generalised trust. However, we manipulated a variety of factors to make one character seem trustworthy and the other untrustworthy, and some of our manipulations might have indicated competence or efficiency (e.g. being on time vs. late, or seeming reliable vs. unreliable) rather than trustworthiness *per se*. Therefore in the third study we aimed to manipulate trustworthiness in a purer way, using the investment game. Participants played the investment game with a fair character who usually returned a profit, and an unfair character who usually returned a loss, then completed the virtual maze task. In this study, we adapted the virtual maze task for use on a standard desktop computer to demonstrate how it may be used in traditional laboratories without the need for virtual reality software or equipment.

2.5.1 Methods

Participants. Twenty-four participants (14 female, M_{age} = 21.4, SD_{age} = 11.8) were recruited through email advertisements to the ICN departmental participant pool. The sample size was based on Study 2, which had sufficient power with a sample of 24. All participants gave written informed consent and received payment of £5 for half an hour, plus a bonus of up to £3.50 depending on how they played the investment game. The study was granted ethical approval by the MoD Research Ethics Committee (Reference Number 564/MODREC/14) and the ICN Ethics Chair (Project ID Number ICN-AH-PWB-3-3-14a).

Investment game manipulation. In this study we used the investment game to manipulate the trustworthiness of two characters called Anne and Beth. We based our manipulation on the procedure used by Franzen et al. (2011). The participant completes

18 trials with one character who plays fairly (usually returns a profit) and 18 trials with one character who plays unfairly (usually returns a loss). Participants played the fair and unfair characters in a counterbalanced order, and the names of the characters were counterbalanced.

At the start of each trial the participant had £1 to play with (this did not accumulate over trials). They could invest any proportion of the £1 with the character (e.g. Anne), who would always triple their investment. Then Anne would choose a proportion of the tripled amount and return it to the participant. Her decision was displayed on the screen. If Anne returned less than 1/3 of the tripled amount, the participant would end the trial with less than their initial £1, thus making a loss. If Anne returned more than 1/3 then the participant would make a profit. At the end of the trial, the participant saw the amount of money Anne returned and the amount of profit or loss made on that trial. The participant was told that at the end of the experiment we would select one of the 18 trials at random and pay them the outcome of that trial. We clarified that the participant could not lose any of their study payment on this game.

Following Franzen et al. (2011), the fair character returned 41.3% of the tripled investment on average. Their returns ranged in nine steps from a maximum return of 66% to a minimum return of 17% (two thirds of trials led to a profit). The unfair character returned 25% of the tripled investment on average. Their returns ranged in nine steps from a maximum return of 50% to a minimum return of 0% (two thirds of trials led to a loss). The trial order for each character was presented in a pseudorandom sequence, which was the same for all participants.

Virtual maze task. We adapted the virtual maze task for use with a standard desktop computer. The participant viewed the task on a standard monitor and made responses using the keyboard. They also wore headphones to hear the advice from each character.

Virtual maze environment. The maze was generated from the same virtual rooms and corridors as in Study 2. However, in Study 3 there were no hologram chambers and the participant could not move around the virtual space. We also changed the blue and red doors to brown so that they would have equal visual value.

Trial procedure. The participant completed 12 trials as in the previous studies. The trial sequence is summarised in Figure 2-4. Each trial began with a video clip where the camera view moved through a corridor and a room until it ended up facing two doors at the far end of the room. At the end of the video clip the view of the doors stayed on the screen. Then a black panel and fixation cross were superimposed in the centre of the screen. After 1.5s, the panel displayed silhouettes and names of each virtual character, along with phone icons. One character (e.g. Anne) always appeared on the left and the other (e.g. Beth) always appeared on the right, in a counterbalanced fashion. The participant had an unlimited amount of time to make a response using the keyboard. If they chose to call a character for advice, that character's phone icon would turn green and her verbal advice was played via headphones, e.g. *'I think it's the left door this time'*. The participant could then call the other character for advice if they chose. The trial ended when the participant chose which door to go through, which triggered a short video clip of the door opening and the camera view moving forward through the door. If the participant chose a door without asking for advice, that trial was recorded but did not count towards the requisite twelve trials.

Anne and Beth were programmed to advise the left door in half the trials and the right door in the other half. They were also programmed to advise the same door as each other in half the trials and different doors in the other half. In order to generate the verbal advice stimuli, we pre-recorded twelve scripted phrases the same as the previous studies.

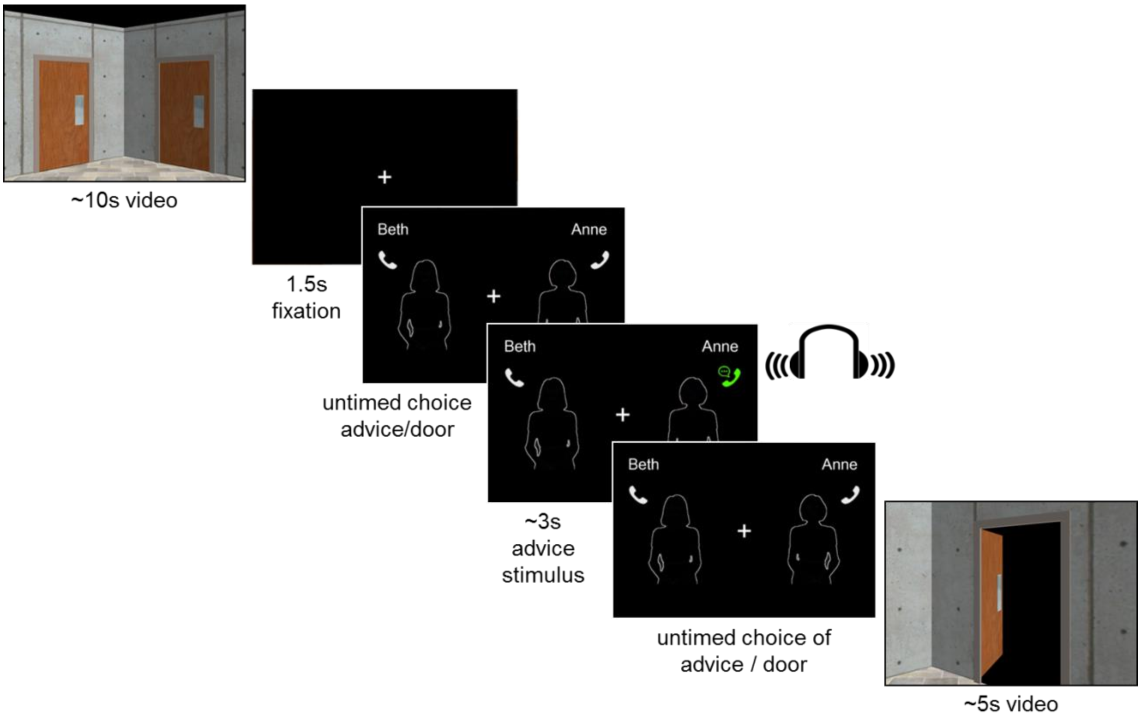

Figure 2-4. Example trial sequence in the maze task adapted for desktop computer (Study 3). The participant sees a video clip where the camera moves through the maze. A black panel with fixation cross is superimposed on the last frame of the video. The participant makes untimed choices to hear advice from Anne or Beth, or choose the left or right door. After choosing a door, the participant sees a video clip where the camera goes through the door.

Procedure. Davide Paoletti carried out the data collection.

Investment game manipulation. First, the experimenter gave a verbal introduction to the investment game and then the participant received instructions on the computer screen. The participant practiced the investment game and was instructed that they were not playing for real money in the practice. They completed four practice trials of the investment game with a character called Kate, who always returned 40% of the tripled investment on each trial (a small profit of up to 20p). When they had completed the practice, the participant was reminded that the rest of the game would be played for real money. Then they completed 18 trials of the investment game with one character (e.g. Anne), followed by 18 trials with the other character (e.g. Beth).

Virtual maze task. After completing the investment game with both characters, the experimenter gave a verbal introduction to the maze game and asked the participant to put on headphones. Then the participant received instructions on the computer screen. First they completed four practice trials of the maze game, in which the four possible advice combinations from Anne and Beth (both say left; both say right; Anne says left but Beth says right; Anne says right but Beth says left) were included in a randomised order. When they had completed the practice, the participant was told that they would start the maze again from a new position in the main task. Then the participant played the maze game until they had completed 12 trials where they asked for advice from at least one character.

Post-study questionnaire. Finally, the participant completed a questionnaire about the study on the computer. First, they were asked about the purpose of the study. Second, they were asked to rate two items about the fairness and trustworthiness of each character as a manipulation check. Fairness was rated was rated on a 7-point scale from '*extremely unfair'* to '*extremely fair'*. Trustworthiness was rated on a 7-point scale from '*I did not trust [Name] at all*' to '*I trusted [Name] completely'*. Lastly, the participant was asked describe any strategy they used in the investment game.

At the very end of the study, participants received payment. The experiment took approximately 30 minutes.

2.5.2 Contributions

Joanna Hale programmed the investment game and the virtual maze task. Audio Stimuli were recorded by Alexandra Georgescu. Davide Paoletti recruited participants and carried out data collection. Joanna Hale carried out the analyses that follow.

2.5.3 Results and Discussion

Manipulation checks. As a manipulation check, we carried out paired sampled *t*-tests to see whether participants' ratings of fairness and trustworthiness differed for the fair and unfair characters. Participants rated the fair character as significantly more fair during the investment game (M = 4.38, SD = 1.58) than the unfair character (M = 3.08, SD = 1.47; $t(1,23)$ = 2.23, p = .04, d = 0.17). They also rated the fair character as significantly more trustworthy (*M* = 4.91, *SD* = 1.59) than the unfair character (*M* = 3.25, *SD* = 1.48; *t*(1,23) = 2.96, *p* = . 007, *d* = 0.60). In addition, participants invested significantly more money with the fair character (*M* = £0.67, *SD* = £0.21) than the unfair character (*M* = £0.38, *SD* = £0.19) (*t*(1,23) = 6.31, *p* < .001, *d* = 1.29). Therefore, the manipulation was successful in making one character seem more trustworthy than the other.

Trust in the Virtual Maze. To test whether participants trusted the fair and unfair characters differently in the virtual maze, we carried out paired-samples *t*-tests on each of our dependent variables from the maze task. Participants followed advice from the fair character more than the unfair character, although this result was marginally significant (*t*(1,23) = 2.03, *p* = .055, *d* = 0.41). The Bayes Factor for this test (BF_{10} = 1.21) indicates anecdotal support for the hypothesis that people follow the advice of the fair character more than the unfair character. There were no significant differences in how often participants called each character overall (*t*(1,23) = 1.51, *p* = 0.15, *d* = 0.31, BF_{10} = 0.58), or called each character first (*t*(1,23) = 0.83, *p* = 0.42, *d* = 0.17, BF_{10} = 0.29).

When we examined participants' responses about any strategy they used in the virtual maze, 10 participants (42%) said that they based their maze decisions on how the characters played the investment game. Ten participants (42%) based their maze decisions on how the characters sounded when giving advice. Three participants (13%) indicated that they used a mixture of these strategies. One participant (4%) reported they had no strategy. Participants who used a strategy based on the investment game followed the advice of the fair character significantly more than the unfair character when they played the virtual maze game (*t*(1,9) = 2.75, *p* = 0.02, *d* = 0.87, BF_{10} = 3.26). The Bayes Factor for this test indicates moderate evidence in favour of the effect, despite the reduced sample size.

Overall, these results suggest that the investment game manipulation was successful but did not lead to strong effects on behaviour in the virtual maze because a large proportion of participants (42%) used audible voice cues to inform their maze decisions instead. If participants had not been distracted by these cues, we would expect to see a stronger effect of fairness on behaviour in the virtual maze. It is important to note that the spoken delivery of the advice stimuli was well-matched for the fair and unfair characters, and stimuli were presented in a randomised order. This means that the participants who relied on any subjective audible cues present in the advice stimuli would have provided very noisy data. However, even with this limitation we were able to detect a marginal difference in how much participants followed the advice of each character in the virtual maze task. We did not see effects on how often each character was called for advice (either called first or called at all), which could be because pressing a button to ‘telephone’ a character is much less socially salient than approaching an embodied virtual character in an immersive 3D environment. Our findings suggest that the simplified format of the task in this study could be less suitable than the full VR version for measuring implicit trust in terms of approach behaviour.

2.6 General Discussion

Across three experiments, we aimed to test whether our novel virtual maze task was sensitive to differences in trustworthiness between two virtual characters. In the first two studies, participants briefly interviewed a trustworthy character (Mike) and an untrustworthy character (Ryan) before completing the virtual maze task, which was either presented on a projector screen (Study 1) or in immersive virtual reality using an HMD (Study 2). In the third study we manipulated the trustworthiness of two characters called Anne and Beth using an investment game and participants completed a low-tech version of the virtual maze task on a standard computer. We measured three outcomes in the virtual maze: (1) the proportion of trials on which each character was approached for advice; (2) the proportion of trials on which each character was the first (or only) character approached for advice; and (3) the proportion of trials on which the participant followed each character’s advice.

In the two VR studies, we found that people approached the trustworthy character significantly more than the untrustworthy character. The trustworthy character was also significantly more likely to be the first or only character they approached on a given trial (Study 2). Furthermore, we found that people followed advice obtained from the trustworthy character significantly more often than advice from the untrustworthy character. This suggested that all three measures we took from the virtual maze were sensitive to our manipulations of trustworthiness, supporting the effectiveness of our paradigm for measuring specific trust. However, in both of the VR studies we manipulated multiple aspects of verbal and nonverbal behaviour to achieve a large effect of trustworthiness, and therefore detected a large effect on each of the virtual maze measures (Table 2-7). In Study 3, we manipulated trustworthiness in a more subtle way by making characters play fairly or unfairly in an investment game (Franzen et al., 2011; Lis et al., 2013) and we implemented a low-tech version of the maze task without VR. Perhaps due to a combination of these design factors, the virtual maze task was not able to detect such strong effects in the third study (Table 2-7).

Table 2-7. Effect sizes.

Measure	**Effect size (*d*)**		
	Study 1	**Study 2**	**Study 3**
Trust rating	3.10	2.11	-
Investment game	0.98	0.94	-
Maze: Approach for advice	1.01	0.99	0.58
Maze: Approach first	0.89	0.97	0.29
Maze: Follow advice	1.63	2.06	0.41

As well as our virtual maze task, in our two VR experiments we also included two of the major existing methods for measuring trust: self-report ratings and the investment game. We found a large significant effect of our trustworthiness manipulation on these measures. In the next sections we discuss what trustworthiness ratings, the investment game and the virtual maze task are actually measuring. In particular, we evaluate the effectiveness of each approach for measuring levels of specific trust towards a stranger.

2.6.1 What do Trustworthiness Ratings Measure?

One of the most straightforward and direct methods to measure trust is through self-report ratings, although they are often criticised for being open to ambiguity (Ben-Ner & Halldorsson, 2010; Hooghe et al., 2009) and social biases (McCambridge et al., 2012), and having a poor correspondence with behaviour (Armitage & Christian, 2003). In this study we asked participants to rate their perceptions of trustworthiness and rapport towards the virtual characters, by indicating their agreement with statements like 'I think Mike is very reliable'. The scales we used to measure trust and rapport have been used in other VR research in our lab and are similar to other validated questionnaires (McCroskey & Teven, 1999). However, in the present study we did not carry out a formal validation of our scales, which limits the conclusions we can draw about their effectiveness. When we tested for the presence of order effects in Study 2, we found that rapport ratings were affected by the order in which participants met the trustworthy and untrustworthy characters. Although this finding does not have a direct bearing on trustworthiness ratings, it highlights a general limitation of this method for repeated measures designs.

2.6.2 What Does the Investment Game Measure?

We suggested in the introduction that how much people invest in the investment game may reflect a mixture of their level of generalised trust and their specific trust towards the other player (trustee). Our data supports this suggestion. We found that participants' investments with Mike and Ryan were highly correlated, which is consistent with previous pilot studies we have conducted. This means that participants who invested a lot with Mike were also likely to invest a lot with Ryan, reflecting a stable individual characteristic. The correlation did not merely reflect similar perceptions of each character's trustworthiness, since trustworthiness ratings were (non-significantly) negatively correlated across characters. Therefore, participants perceived a difference in the characters' trustworthiness, but nevertheless invested similar amounts with each one. We interpret this as evidence of generalised trust, although previous research indicated that other stable individual traits such as risk-aversion or altruism might possibly have determined this outcome (Ashraf et al., 2006; Ben-Ner & Halldorsson, 2010; Cox, 2004; Karlan, 2005; McEvily et al., 2012; Schechter, 2007).

Our data from Study 2 reveal further how generalised trust and specific trust may each play a role in the investment game. We found a significant order effect showing that the ratio of participants' investments towards the two characters depended on whether they interacted with the trustworthy character first, or the untrustworthy character. From decomposing this effect it was clear that participants invested similar amounts of money with the first character they met, regardless of whether it was the trustworthy or the untrustworthy character. In other words, when they had no 'baseline' comparison from interacting previously with a virtual character, we assume that participants invested money based on their level of generalised trust. In contrast, people invested much more with the trustworthy character than the untrustworthy character once they had already completed one interaction. The first interaction may have provided a reference point so that when they played the investment game for a second time, participants made decisions based on specific trust towards the other player. This may be a limitation in studies using repeated investment games to measure trust towards different individuals.

2.6.3 What Does the Virtual Maze Measure?

Our aim in designing the virtual maze task was to develop a new way of measuring specific trust towards a stranger. Our results indicate that the virtual maze task was successful in capturing specific trust rather than generalised trust. In Studies 1 and 2, we did not find a significant relationship between how often participants approached one character and how often they approached the other character, which suggests that people did not simply act according to their own level of generalised trust in others. If the task were mainly measuring generalised trust, we would expect participants to approach both characters a similar amount overall. The other two measures we took from the virtual maze are actually anti-correlated across characters due to the nature of the task: e.g. if someone approached Mike first they could not also approach Ryan first, and on 50% of trials if they followed Mike's advice they could not also follow Ryan's (on the other 50% of trials the advice was the same).

These forced-choice aspects of the virtual maze task may make it particularly sensitive to specific trust rather than generalised trust, because the participant is implicitly forced to assess the trustworthiness of one character versus the other, much like in the ask-endorse paradigm when the participant has to select one informant versus the other. We would expect this forced-choice approach to be much more sensitive than asking the participant to make a choice about one character alone. Indeed, when participants had to choose which character's advice to follow in the VR maze (studies 1 and 2) we found a larger difference between Mike and Ryan compared to when participants had to choose how much money they would invest with each character individually (Table 2-7). However, our results from Study 3 suggest that outside of VR the virtual maze task may be less sensitive to these effects, particularly if other cues are present which might affect participants' decisions, such as tone of voice.

In Study 3, we addressed the possibility that our manipulations of Mike and Ryan might have included elements of competency or effectiveness rather than trustworthiness. According to Fiske et al.'s model of social cognition (Fiske et al., 2002, 1999), warmth (which includes perceived trustworthiness) is evaluated independently of competence. In addition, economic definitions suggest that trusting involves risking exploitation, which was not possible in our virtual maze task (i.e. the virtual characters could not exploit the participant for gain). Therefore, it could be argued that our results from Studies 1 and 2 reflected perceptions of Mike and Ryan's competence or reliability and that this does represent trust *per se*. Therefore, in Study 3 we manipulated trustworthiness of two female characters (Anne and Beth) using the investment game instead. We also measured trust towards each character using a low-tech version of the maze task. In this version, participants could ask Anne and Beth for advice in a ‘phone a friend’ fashion by pressing a call button. Although we found weaker effects using this simplified version of the task, 42% of participants explicitly reported that they made decisions in the maze task based on how the characters behaved in the investment game and these participants followed the advice of the fair character significantly more than the unfair character. Therefore, Study 3 demonstrates that the maze task is able to detect differences in trust that are based on previous economic exchange rather than impressions of competency. We did not directly compare trustworthiness and competency manipulations in any of our studies, so it is hard to say to what extent each factor might determine behaviour in the virtual maze; however, previous research using selective trust paradigms with children has shown that benevolence is favoured more than competence or expertise when selecting an informant to trust (Johnston et al., 2015; Landrum et al., 2013). We should also note that 42% of participants made decisions in the maze based on how confident each character sounded when delivering verbal advice. This was an unintended outcome that added noise to our data, but it also highlights a possibility for future research. By manipulating social cues *within* the virtual maze it could be possible to investigate the interaction of these cues with pre-formed impressions of trustworthiness.

Finally, although we designed the virtual maze task to measure specific trust in one person, it would also be possible to derive a measure of generalised trust from the task. This could be calculated as the number of trials where the participant asks anyone for advice, rather than simply choosing an option on their own (in our studies such trials were not counted). A next step would then be to see how the virtual maze task compares to other existing measures of generalised trust. Unfortunately, the initial studies we present here did not include large enough samples to robustly test the relationship between the virtual maze task and existing questionnaire or economic measures of trust. To see how the virtual maze task correlates with other measures, it will be necessary to conduct a larger correlational study with sufficient power.

2.6.4 Methodological Advantages of the Virtual Maze

The virtual maze has several methodological advantages that may make it a valuable tool for future studies. Firstly, the task offers high ecological validity in the form of a fairly realistic scenario. While we may not find ourselves in mazes during everyday life, we often have to make choices about whether to ask a stranger for advice and follow what they say. In such situations our trust is implicit, and this is reflected in the virtual maze task. Unlike ask-endorse paradigms used with children (Clément et al., 2004; Koenig & Harris, 2005), we did not have to explicitly ask participants who they wanted to approach in the maze and whether they would follow that person's advice. These possibilities were obvious from people's everyday experiences of navigating an unfamiliar place. Similar to the ask-endorse paradigm, the maze task is also designed for making comparisons across multiple targets and allows us to average trust behaviour over multiple trials. Finally, the virtual maze task offers many opportunities to measure other implicit aspects of trust behaviour, such as the speed at which people approach different virtual characters, the time they spend looking at each character, or how direct a pathway they take through the rooms in the virtual maze.

We have demonstrated that the virtual maze task can also be adapted for a traditional computer display, avoiding the need for virtual reality software. Future studies that use this version of the task should consider replacing the verbal advice stimuli with simpler stimuli (e.g. an arrow or word 'left') so as to avoid confounding effects of confidence perceptions based on the verbal cues. In addition, to get strong effects it could be important to increase the social saliency and effort involved in calling one character or the other. Despite these limitations, adapting the task for a traditional computer display had the advantage over the VR versions that it avoided nausea associated with navigating through a virtual space (Davis, Nesbitt, & Nalivaiko, 2014; LaViola, 2000). Although we attempted to reduce these effects by incorporating smoother movement and fewer twisting routes when we adapted the VR task for HMD (Study 2), several participants were still unable to complete the task. A virtual reality platform where participants can see their physical body moving, such as the CAVE (Cruz-Neira, Sandin, & DeFanti, 1993), may provide a more suitable way to administer the maze task in virtual reality.

2.7 Conclusions

In this chapter we have described a novel task for measuring specific trust towards a stranger. We showed that our virtual maze task is able to detect differences in specific trust towards different target individuals. Importantly, we found that behaviour in the virtual maze task appears to reflect specific levels of trust towards each target rather than a generalised level of trust, which was our aim in designing the task. In contrast, we found that behaviour in the investment game reflected generalised trust as well as specific trust. The virtual maze task also has several practical advantages over self-report ratings and existing behavioural measures. In particular, it involves more implicit trust decisions than the investment game. Administering the task in virtual reality may provide opportunities to measure other implicit behaviour, such as proximity to the target individuals. However, the virtual maze task can also be administered on traditional desktop displays, although this may yield weaker effects. Overall, the virtual maze task could provide an alternative behavioural method for researchers interested in factors that affect specific trust towards unknown or familiar individuals. Future steps will be to establish how behaviour in the virtual maze correlates with other measures of trust, and use this new tool to explore other factors, such as mimicry, which may affect people's specific trust towards others.

In Chapter 4, we go on to use the maze task to test the effect of mimicry on trust. We did not use the maze task in our first mimicry study (Chapter 3) because that study was completed before we developed the maze task.

2.8 Appendix

Study 1 Scripts

Interviewer	Mike	Ryan
	Hi, my name's Mike. Nice to meet you! I'm really looking forward to doing this experiment with you! But I think you've got some questions for me first, right?	Hey, I'm Ryan. Umm... so, yeah... I have to answer some questions or something. So, shall we get on with it?
What is your occupation?	Well, I'm an engineering student at UCL. I'm in third year of my PhD now, and it's all about developing safety features for aircrafts. Before I started, I was working for a specialist aircraft manufacturer as a quality engineer, and that got me really interested into, umm, new safety features. Part of my project involves collaborating with my old colleagues to test out new prototypes – it's been really good to work with them. When I finish my PhD I'm hoping to carry on working with the same company.	Er... it's kind of complicated to explain. I'm a junior associate consultant with my Uncle's finance firm. Um... basically I deal with acquisitions and I assess the NAV of the company clients and review audits. It's, er... it's probably a bit technical for you to understand. I was a bit surprised I got hired, actually, because, er... things didn't go so well at my last job and I basically ended up getting fired after missing too many deadlines. But I've skipped a couple of meetings at this company, and no-one said anything about it, so I guess they're a bit more laid back.
How did you spend last weekend?	Oh, I was pretty busy last weekend actually. On Saturday I did this charity bike ride from London to Brighton – it took most of the day, but it was really worth it: I promised to raise £800 in sponsorship and I managed to smash the target. I was pretty knackered on Sunday, but I didn't get much rest because I had to spend the rest of the day cleaning the flat. Then in the evening one of my friends came round and, umm, we cooked dinner. Luckily I managed to get an early night though, because I had to be at a meeting at 9am on Monday.	Umm, so, last weekend was pretty epic. On Friday night, I took some clients to the Shard and we spent, like, £800 on dinner. And, umm, then we went to this club in Mayfair - I was supposed to help my housemate move his stuff the next day, but there was no way I was getting up in the morning. I, er, feel pretty bad about that actually... but, err... what can you do? So, umm... oh yeah, on Sunday I forgot I was meant to be preparing this client presentation for work - my friend had got us tickets to the rugby and I really didn't want to miss the match, so I called up one of my colleagues and he agreed to do the presentation for me this week. So, er... yeah, that was a relief!
What are your plans for the summer?	My plans for the summer? I've got a couple of different things, really. In July I'm going to Peru with an old school friend. We like to get stuck in to the local culture, so we don't really go to all the touristy places. I can't wait to get on the plane to be honest! I've also got an internship starting at an aeronautics lab – so that's when I get back from South America. I guess I've got a pretty busy summer, so I'm working quite hard at the moment to try and get all my work finished before then.	Umm... I guess I'll probably head off somewhere exotic for a few weeks– maybe, like, a cruise or something. You know, lie in the sun and drink cocktails... that's pretty much all I want to do. Maybe Dubai – our company has some big clients there, so, err, I might be able to get them to pay for my flights out there and then just, like, have a bit of a holiday. Or it would be cool to go to Vegas or something and play all the big casinos. I think here's some company teamwork retreat happening as well, but I'm going to try and get out of that. So, err, yeah, that's pretty much it.

Study 2 Scripts

Interviewer	Mike	Ryan
	Ok, yep, I'm ready to start! Hey I'm Michael, but everyone calls me Mike! I'm going to be doing this study with you today, I'm looking forwards to it - shall we get started?	Ahh sorry I'm late yeah ok I'm ready now! Hey, I'm Ryan and I'm going to be doing this study with you today. I'm looking forwards to it - shall we get started?
What do you do for a living?	Well, I'm a biomedical student in the third year of my PhD now. It's all about developing interventions in surgery to reduce disease. Before I started, I was working as a haematologist at St Thomas' Hospital and that got me really interested in how infections can be controlled in hospitals. I'm really enjoying it, mainly because it involves collaborating across disciplines to test out various antibiotics- it's been really good working with such a range of people.	Well, at the moment I'm working with my Uncle at his restaurant – I'm a waiter. Yeah it's good because the food is really tasty and my uncle is super laid back so he lets me off when I'm late sometimes. I've been doing that for a few years now. I also get along with everyone who works there pretty well, it's kind of like they've become family too, is that really cheesy? Anyway, I really enjoy the teamwork aspect of the job.
What did you do at university?	I studied biomedical sciences in Sheffield. I really enjoyed my time there, I met so many people that I'm still really close with now. I was definitely the nerd of the group though, I ended up graduating with the highest grade in our year.	I studied history, down at Exeter university. Uhh… I didn't think much of the course but the people I met there I know I'll be friends with for a long time. Unis just so relaxed isn't it? I only had a few lectures a week and I used to miss them all the time from being hungover and stuff but no one says anything, well not to my face anyway.
Were you in any university societies?	Yes I was actually, I was a member of the hockey soc all through uni and in the final year I was voted as chair of the society. So I was involved in organising all the matches and training sessions and all of that stuff for a year. Yeah I really enjoyed it, I think it's always a nice feeling to be a member of a team.	Societies? Uhh… Oh yeah, I was part of the film soc. But basically just because they had the best parties. They used to do these big pub-crawls down this one road, and everyone has to have a pint at each one. Fancy dress obligatory of course! Yeah, that was a laugh.
Do you have any hobbies?	I still play hockey just with a local team now. We're a bit better than the uni team was, in fact, we might be competing in a small tournament soon. I think we have a pretty good chance but I don't want to be too confident, we'll have to see on the day.	Umm.. Err… I really just enjoy socialising with friends, maybe playing a couple of video games with some pizza, you know, that kinda thing.
What do you do to relax?	Umm… Well… I like running, I think it's really therapeutic and I like to keep fit and healthy. I'm going to attempt the London half marathon this year and I'm hoping I can maybe raise some money for a nearby homesless shelter. I set up a justgiving page last week and I've already got a few generous donations; now, I just need to make sure I stick to the training schedule I've set up for myself.	Umm… Well… I enjoy cooking, well sometimes I like cooking, me and my flatmates take it in turns to cook meals each day a week, but I kind of messed it up this week, I made this huge pasta bake and forgot all about it, by the time I'd gotten to it it was just a big black mush – totally inedible! It was so annoying! But anyway, I bought everyone a takeaway to make up for it, which is probably nicer than a pasta bake anyway I guess!

How did you spend last weekend?	Last weekend? Umm… oh yeah, it was my mum's birthday on Saturday, she turned 50 so I arranged for some family and friends to surprise her at her local pub, yeah, it was great! I told her it was just me and her going and she didn't suspect a thing. Everyone was hiding behind the chairs and tables in the pub. You should have seen her face when they all jumped out! Priceless!	Last weekend? Umm… oh yeah, god, that was my girlfriend's birthday. I can't believe I forgot about it. I felt awful, it just totally slipped my mind and she definitely clocked on to it. Yeah, she was quite angry. I went over later that day with a huge bunch of flowers to make it up to her, her favourite ones. She forgave me in the end so it's all good, but I certainly won't be forgetting that again!
What are your plans for the Summer?	Ah I can't wait for the Summer! I'll be finishing my PhD then, and I'm going to take a little holiday and go away travelling with an old friend of mine. We just booked our flights the other day actually, we're flying into Bangkok and then out of Hanoi, in Vietnam a month later. So we're going to travel between the two. I'm really excited, it'll be great to really get into the culture and try something completely different.	Umm… Well I haven't really made that many plans yet, I should probably get started on that now you mention it… I think it'd be good to do an internship or something, maybe related to teaching if I could find one, I'm sure they have like websites for that kinda thing though, right? If I can't find anything I'll just go on holiday with my family. That way my mum can do all the organising and I can just tag on at the end. But yeah, other than that I should probably look into some internships soon…

Chapter 3. Using virtual reality to test whether mimicry leads to trust and rapport

3.1 Abstract

People mimic each other's actions and postures during everyday interactions. It is widely believed this mimicry acts as 'social glue', leading to positive social outcomes such as increased rapport towards the mimicker. In this study we develop a new mimicry paradigm using virtual reality to provide a stricter test of this claim and to explore whether the precise timing of mimicry affects how people respond to being mimicked. Fifty participants interacted with two virtual characters who either mimicked their head and torso movements at a 1 or 3 second time delay, or did not mimic. Participants rated the smoothness of the interaction and their feelings of similarity, rapport and trust toward each virtual character. Rapport was higher towards the virtual character that mimicked, in line with the social glue theory, and this was not affected by the timing. There were no other significant effects of mimicry or mimicry-timing interactions, although the paradigm achieved a high level of social realism. These findings suggest that virtual mimicry is sufficiently realistic to generate positive social effects, but controlling for all other social signals may lead to smaller mimicry effects than traditionally expected.

3.2 Introduction

As we outlined in Chapter 1, it is widely believed that mimicking another person has positive consequences for the social interaction (Chartrand & Lakin, 2013; Lakin & Chartrand, 2003; van Baaren, Decety, et al., 2009; van Baaren et al., 2004). Based on this belief, the prominent 'social glue' theory of mimicry proposes that mimicry is an adaptive social strategy that helps people to bond together by inducing feelings of connectedness and liking (Lakin et al., 2003). In line with the prevailing view that mimicry leads to liking, affiliation and rapport, mimicking others has been advocated as a strategy for business and personal interactions, as well as teaching and therapy (Bernieri, 1988;

LaFrance & Broadbent, 1976; Scheflen, 1964). However, while there is strong evidence that wanting to affiliate with someone makes you more likely to unconsciously mimic them, the evidence that people actually like mimickers is less robust (see Chapter 1; Hale & Hamilton, 2016a). In addition, the cognitive processes involved in responding to being mimicked are currently unclear. In particular, we do not know how mimicry is unconcsciously 'detected' by the mimickee, and it is unclear whether positive feelings towards the mimicker depend on detecting general contingency (Catmur & Heyes, 2013) or predictability (Friston, Mattout, & Kilner, 2011; Kilner, Friston, & Frith, 2007) in their behaviour, or whether we have a cognitive mechanism that responds specifically to mimicry (Hale & Hamilton, 2016a).

For all of these reasons, it is important to probe how people respond to being mimicked in detail and with precise experimental control. Many previous studies of this topic have used naturalistic situations in which a confederate is trained to subtly mirror the movements and posture of a naïve participant during a social interaction task such as describing photos to each other (Refer to Chapter 1; Chartrand & Bargh, 1999; Stel et al., 2011; van Baaren et al., 2004). Headline results from these paradigms suggest that mimicry increases prosocial behaviour (Ashton-James et al., 2007; Fischer-Lokou et al., 2011; van Baaren et al., 2004) and may lead to a greater feeling of closeness (Van Baaren & Chartrand, 2005), liking (Kouzakova, van Baaren, et al., 2010; Stel et al., 2011) and trust towards the mimicker (Maddux et al., 2008), although these results may be modulated by the social context and individual characteristics (Dalton et al., 2010; Leander et al., 2012; Stel et al., 2011). However, there are several reasons to be cautious of accepting these naturalistic studies of mimicry at face value. First, both effect sizes and experimental power have been small in many previous studies, and there may be false-positives present in this literature (Chapter 1). Second, it is not always clear whether a confederate can accurately change her mimicry behaviour without also changing other behaviours which may naturally vary alongside mimicry, such as eye contact, smiling, or the level of contingency between the mimicker and mimickee. Third, confederate behaviour may be implicitly affected by knowledge of the experimental condition they have to perform, or the cognitive demand of the instruction to mimic or avoid mimicking (Stel et al., 2009). Finally, even well trained confederates lack control over the exact timing and matching precision of their movements, making it hard to draw conclusions about what it is that makes people unconsciously detect when they are being mimicked and change their evaluation of the mimicker.

A strong test of the claim that mimicry itself leads to positive social consequences can come from virtual mimicry paradigms in which every parameter of the mimicry interaction is precisely controlled. Given that people react towards virtual characters similarly to real people (Bailenson et al., 2001; Garau et al., 2005), virtual reality is becoming an increasingly popular tool in social psychology and neuroscience. Bailenson and Yee (2005) generated mimicry in virtual reality by tracking participant's head movement and applying the same movement to a virtual character after a delay. In the control condition, movements from the previous participant were applied to the virtual character instead. Bailenson & Yee (2005) found that, after listening to a persuasive speech from the virtual character, participants evaluated it more positively if they had experienced mimicry rather than pre-recorded motion. Other virtual mimicry studies also suggest virtual characters can generate similar mimicry effects to human confederates (Hasler et al., 2014; Verberne et al., 2015; Vrijsen, Lange, Dotsch, et al., 2010). However, this approach is still relatively new and results have not always been replicated, even within the same research groups repeating the same computer code (Verberne et al., 2013). Therefore, establishing if being mimicked by virtual characters does lead to increased rapport and trust would provide a strong test of the social glue hypothesis and an important tool for studying the cognitive mechanisms which detect when someone is being mimicked.

One aim of the present study was to build a new virtual mimicry paradigm and test whether this would replicate the positive consequences of mimicry found in studies using human confederates. Our virtual mimicry paradigm builds on Bailenson & Yee's

(2005) algorithm for mirroring motion-tracked movements. However, we incorporated this into a more socially interactive task in which the participant and virtual character take turns to describe photos to each other, which has commonly been used in confederate studies (Chartrand & Bargh, 1999) and may provide a closer comparison with this body of work. In this initial study we aimed to explore a range of possible mimicry outcomes, and therefore we tested several previously reported outcomes, including self-other overlap (Van Baaren & Chartrand, 2005), interactions smoothness (Chartrand & Bargh, 1999) and feelings of similarity, rapport (Bailenson & Yee, 2005; Chartrand & Bargh, 1999; Kot & Kulesza, 2016; Lakin & Chartrand, 2003) and trust (Maddux et al., 2008; Verberne et al., 2013, 2015). A second aim of this initial exploratory study was to provide a proof-of-concept that virtual reality can be used to systematically test mimicry parameters that have previously been difficult or impossible for human confederates to control reliably. One such parameter which was flagged in a recent review (Chartrand & Lakin, 2013) is the precise timing of the mimicked movements.

3.2.1 The Timing of Mimicry

If there is a cognitive mechanism which responds when another person is mimicking me, we would expect it to have some window tuned to the timing of mimicry, whereby mimicry at short delays is easier to detect or has stronger consequences than mimicry at longer delays. Such a window might also be tuned to the natural timing of mimicry found in social interactions. However, there is little data about the role of this factor or the natural timing of spontaneous mimicry (Chartrand & Lakin, 2013). Some claim that mimicry naturally occurs within a 2-5s time window (Leander et al., 2012; van Baaren, Decety, et al., 2009), and this timescale has sometimes been used when training confederates (Leander et al., 2012). However, others have implicitly adopted different timescales, such as 1-2s (Kot & Kulesza, 2016; Tanner, 2008), with the most extreme ranging from zero (Chartrand & Bargh, 1999) to ten seconds (Stel et al., 2009), and many studies do not report a timescale for mimicry. Therefore, it is unclear whether the timing may impact people’s responses to being mimicked. Preliminary evidence suggests people find it easier to deliberately detect virtual mimicry at a one-second time delay, compared to two, four or eight seconds (Bailenson et al., 2004). However, no studies have systematically tested whether the timing of mimicry affects downstream consequences such as rapport and trust.

We can begin to generate predictions about the role of timing by drawing from literature on contingency in social interactions. When one person mimics another, the mimicker's actions are both similar to and contingent upon what the mimickee just did. The shorter the time delay, the easier it is to recognise the contingent nature of the mimicker's actions (Bailenson et al., 2004). Recognising contingency between our own actions and those of others is intrinsically rewarding and motivating. From infancy, contingent caregiver behaviour increases positive affect, self-efficacy and social motivation towards the caregiver (Dunham, Dunham, Hurshman, & Alexander, 1989; Millar, 1988; Watson & Ramey, 1972). Therefore it has been suggested that contingency, rather than similarity, may be the 'active ingredient' in mimicry that produces positive responses (Catmur & Heyes, 2013). This suggestion was supported by the finding that people respond positively to contingent movements regardless of how similar the movements are to their own (Catmur & Heyes, 2013). Other research shows that being in synchrony has similar positive effects to being mimicked, including liking (Hove & Risen, 2009; Miles, Nind, & Macrae, 2009) and prosocial behaviour (Reddish, Fischer, & Bulbulia, 2013; Piercarlo Valdesolo, Ouyang, & DeSteno, 2010; Wiltermuth & Heath, 2009), and that synchrony and mimicry activate similar reward regions in the brain (Cacioppo et al., 2014; Kokal, Engel, Kirschner, & Keysers, 2011). Since synchronised movements are characterised by temporal contingency rather than similarity, this evidence is also consistent with the idea that positive mimicry effects may depend on contingency. Based on these observations, we would tentatively predict that people will have more strongly positive responses to mimicry with shorter time delays, when the contingency in the mimicker's actions might be more salient.

3.2.2 The Present Study

Virtual mimicry is an ideal method for exploring the question of timing, since virtual characters can be programmed to repeat the participant's actions after a precise, constant time delay. Building on Bailenson and Yee's (2005) approach, we implemented a virtual reality version of an interactive photo description task commonly used in confederate studies (Chartrand & Bargh, 1999; van Baaren, Decety, et al., 2009). In this task, the participant and an virtual character take turns to describe photographs to each other, giving the feeling of interaction without a full conversation (Figure 3-1). We used two virtual characters: one mirrored the head and torso movements of the participants after a specific delay (1 second or 3 seconds); the other virtual character showed pre-recorded natural head and torso movements without mimicry. Participants interacted with the mimicking and non-mimicking virtual characters one after the other, in a within subjects design. The time-delay of mimicry was a between-subjects factor. After interacting with each virtual character, participants completed a number of ratings to evaluate their feelings about that virtual character, including rapport, trust, similarity, the smoothness of the interaction and self-other overlap (Aron et al., 1992), i.e. feelings of closeness towards others. A co-presence questionnaire was used to evaluate the realism of the VR (Friedman et al., 2014; Pan et al., 2012). At the end of the study, participants were carefully debriefed to determine if they consciously detected mimicry.

3.3 Methods

3.3.1 Participants

A power calculation using G*Power (version 3.1.7, 2013) showed that 46 participants would be sufficient to detect an effect size of η_p^2 = .07 (a conservative estimate based on previous studies) with power of 0.95. Therefore recruitment continued until we obtained 50 participants who did not consciously detect the mimicry manipulation. Altogether, sixty three participants (44 female, M_{age} = 26) were recruited through email advertisements. All participants gave written informed consent and received £7.50 payment for 1 hour. There were a total of 26 participants in the 3s group and 37 participants in the 1s group. Ethical approval was obtained from the UCL Institute of Cognitive Neuroscience Ethics Chair (Project ID Number ICN-AH-PWB-3-3-14a) and the Ministry of Defence Research Ethics Committee (501/MODREC/13). The methods were carried out in accordance with the approved guidelines and the Declaration of Helsinki (2013). All participants gave written informed consent to take part.

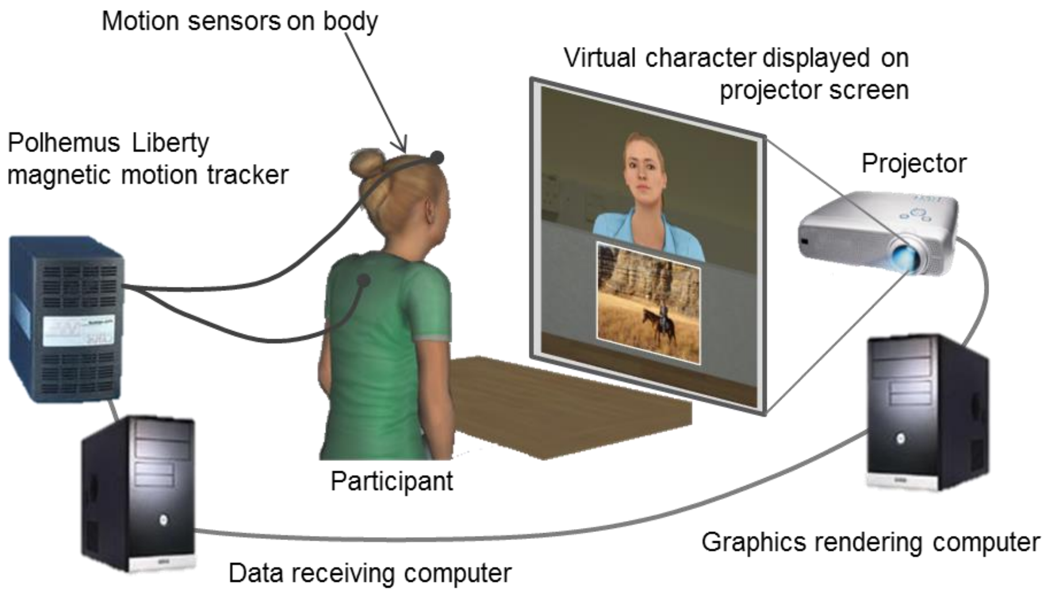

Figure 3-1. Virtual mimicry system. The participant was fitted with two motion tracking sensors which recorded their head and torso movements. The movement was mirrored by a virtual character after a specified time delay. The virtual character was displayed on a projector screen.

3.3.2 Equipment

Participants sat at a desk facing a 90 x 160cm projector screen at approximately 1m distance (Figure 3-1). All parts of the experiment were presented on the projector screen. The participant used a mouse to make responses.

We used a Polhemus magnetic motion tracking device (Polhemus Inc., Vermont) to record participants' movements. This device detects the position and rotation of sensors within a low-intensity electromagnetic field, with a specified frequency (we used

120 Hz). Participants were fitted with two sensors. One sensor was fixed on the participant’s forehead using a Velcro cloth band. Another sensor was attached to the participant’s upper back using microporous tape.

3.3.3 Virtual Reality Environment

We used Vizard virtual reality software (Worldviz, 2014) to display virtual characters in a virtual environment that looked like an extension of the physical laboratory. The physical walls and the participant’s desk appeared to extend into the virtual space, and the virtual characters appeared as life-sized people seated on the other side of the desk, facing the participant (Figure 3-1).

3.3.4 Virtual Mimicry Algorithm

To achieve the mimicry condition, we used two separate computers networked together (Figure 3-1). One computer received data from the Polhemus motion tracker. The data was applied onto a digital mannequin in MotionBuilder (Autodesk, 2014). Inverse kinematics were applied so that the whole body of the mannequin would move in a biologically realistic way based on limited data points (i.e. two motion sensors). While MotionBuilder was running on one computer, we used a software plugin called Live Characters (Worldviz, 2014) to stream the mannequin’s movement over to another computer running Vizard (WorldViz, 2014). In Vizard, we programmed a virtual character to mirror the movement of the mannequin after a delay of 3s or 1s, depending on the experimental condition. In addition, a restriction was imposed so the pitch (up/down tilt) of the participant’s head would only be mimicked within a range of 0 to -10 degrees from neutral. This restriction was designed to avoid mimicry of upward head tilts, which appeared unnatural in a pilot experiment. The virtual character from Vizard was displayed to the participant on the projector screen. Thus, the participant saw a virtual character which mirrored their movements in a biologically realistic way after a short delay.

For the non-mimicry condition, we animated a virtual character using pre-recorded movement from a pilot participant who was experiencing the mimicry condition. The pilot participant moved a moderate amount while being mimicked and did not detect mimicry. We used this pre-recording for all participants, instead of ‘yoking’ each participant to the previous one (Bailenson & Yee, 2005; Vrijsen, Lange, Dotsch, et al., 2010), so that the non-mimicry condition would be the same for all participants.

3.3.5 Virtual Characters

We prepared three virtual characters for the experiment. The first virtual character (a male named Mike) was only used for practising the experimental tasks. The main two virtual characters in the experiment were females named Anna and Becky. The virtual characters were selected from the ‘Complete Characters’ set provided with Vizard software. Anna and Becky’s appearances were selected according to two criteria: firstly, that they appear similar to the average age of our recruitment population (university students aged 18-30 years, with mean sample age of 26 years); and secondly, that they appear similar to each other, e.g. both having blonde hair and casual clothing. This was because piloting with a different Becky character, who had brown hair, showed that participants significantly preferred the blonde-haired Anna to brown-haired Becky, and we wanted to avoid any such differences that might obscure mimicry effects.

Apart from mimicry behaviour, we also programmed each virtual character to perform some other social behaviours when they interacted with participants. Each virtual character was animated to smile and greet the participant, e.g. *‘Hi, I’m Anna – nice to meet you’*. Whenever the character spoke, they moved their jaw according to the amplitude of a pre-recorded audio file. If the character was describing a photo, they alternated between tipping their head to look down at the photo they were describing (for 5-8s), and looking up at the participant with direct gaze (for 1-3s) throughout their turn. If it was the participant’s turn speaking, the virtual character smiled once and then alternated between direct gaze (for 4-7s) and averted gaze (for 0.75-1.5s) throughout the participant’s turn. The virtual character was programmed to blink every 2-5s.

3.3.6 Photo Stimuli and Virtual Character Descriptions

Ten photo stimuli were taken from the National Geographic website, under a category called 'People and Culture', in line with previous mimicry studies (e.g. Chartrand & Bargh, 1999; Cheng & Chartrand, 2003). We pre-recorded Anna and Becky's descriptions of the photographs. Two volunteers with native British accents came to the lab and their voices were recorded using headset microphones. The volunteers were asked to take turns describing the set of 10 photos to each other. They were told to speak for at least 30s per photo (timed by a researcher on a stopwatch) to ensure we obtained at least 30s of recorded description. We later edited each recording to exactly 30s in duration, so that each virtual character would speak for their full turn in the experiment. First, the volunteers practiced describing the photos to each other. Then we recorded them while they described the same set of stimuli to each other (i.e. each volunteer described all 10 photos). During the experiment, the stimuli were randomly divided between Anna and Becky, so that each participant heard a random set of five descriptions from each character.

3.3.7 Experiment Procedure

On arriving in the lab, participants were told they were taking part in a virtual communication study and completed a consent form. They sat at a desk in front of a projector screen and Polhemus motion tracking sensors were fixed to their head and back. A short calibration was performed to allow the VR software to map the participant's motion onto one of the virtual characters. Prior to the study and during the calibration, the participant was simply told that the motion tracker would allow us to record their motion during the experiment. After calibration, the study began.

The participant practiced the task sequence detailed below. Then they completed the task sequence with Anna and then Becky. Either Anna or Becky mimicked the participant (mimicry condition), while the other virtual character did not mimic (non-mimicry condition), in a counterbalanced fashion.

Photo description task (mimicry manipulation). First, the participant completed the photo description task. During the task, the virtual character and participant would take 5 turns each at describing a photo to the other for 30 seconds. The virtual character greeted the participant and smiled at the start, then began the task by describing a photo the participant could not see. At the end of their description, the character smiled again and then it was the participant's turn to describe a photo, which was displayed on a virtual barrier between the participant and virtual character (Figure 3-1). A timer on the virtual barrier counted down from 30s and an audible beep signalled the end of each turn. At the end of the task, the virtual barrier was raised up to occlude the character.

Questionnaires and ratings.

Self-other overlap. Following the photo description task, the participant rated their feelings of overlap towards other people, using a scale based on the Inclusion of Other in the Self scale (IOS; Aron et al., 1992). The participant saw two circles and were instructed to slide the circles further apart or closer together (Verberne et al., 2015). The circles were equally sized and could be moved from 1 diameter apart to completely overlapping. The participant rated their overlap with their best friend (used as a reference point), the virtual character from the photo task, virtual characters in general, and other people in general.

Virtual character ratings. Next, the participant rated the smoothness of the interaction and their feelings of rapport, trust and similarity on a sliding scale from *'strongly disagree'* (scored 0) to *'strongly agree'* (scored 1). The items are given in Table 1. The order of the items within each category (e.g. rapport) was randomised for each participant.

Table 3-1. Items for virtual character ratings.

Construct	Items
Smoothness	I think the interaction with [Character] was smooth. *I think the interaction with [Character] was awkward.
Rapport	I think [Character] is very likeable. I think [Character] is very engaging. I think [Character] is very kind. *I think [Character] is very unfriendly. *I think [Character] is very unpleasant.
Trust	I think [Character] is very trustworthy. I think [Character] is very honest. I think [Character] is very responsible. *I think [Character] is very unreliable. *I think [Character] is very insincere.
Similarity	I think [Character] is very similar to me. *I think [Character is very different from me.

Note. Items marked * were reverse scored.

Co-presence questionnaire. At the end of the experiment, the participant completed a short questionnaire about the level of co-presence they experienced, which refers to the feeling that the virtual characters were really present in the same place as the participant (Garau et al., 2005; Pan et al., 2012). The questionnaire contained four items, structured *'How much did you find yourself reacting to the virtual characters as real people, concerning your...'* (1) thoughts; (2) feelings and emotions; (3) physical responses (4) physiological responses. Participants rated their responses on a 7-point scale from 'not at all' to 'very much so'.

Debriefing. The participant provided written and verbal answers to four questions to determine whether they had noticed the mimicry manipulation or guessed the purpose of the experiment: (1) *'What did you think was the purpose of this study?'* (2) *'Why do you think you were asked to wear the motion tracker?'* (3) *'Did you notice anything unusual about how the virtual characters moved?'* (4) '*Did you notice any differences in how the virtual characters moved?'*

3.3.8 Contributions

Joanna Hale implemented the virtual mimicry algorithm in Vizard (this does not include inverse kinematic modelling which was done automatically in MotionBuilder), recorded all audio stimuli and programmed the photo description task. Joanna Hale carried out all participant recruitment, data collection and analyses.

3.4 Results

Data were excluded from one participant in the 1s group due to technical failure of the motion tracker.

3.4.1 Mimicry Detection

In the 3s mimicry group, one participant out of 26 (3.8%) detected the mimicry manipulation. In the 1s group, 11 out of 36 (30.5%) detected mimicry. The detection rate was significantly greater in the 1s group than the 3s group ($\chi^2(1) = 6.9$, $p < 0.01$). All detectors were excluded from the analyses. Thus, the remaining sample was $N = 25$ in each group.

3.4.2 Co-Presence

We averaged the four items from the co-presence questionnaire into one co-presence score (Cronbach's alpha = .72). The median social presence score was 4.45 (*M*=4.43, *SD* =1.15, *Range* = [1.75, 6.75]) on a scale from 1 to 7, which is very similar to levels found in more immersive virtual reality experiments using a comparable rating scale (Friedman et al., 2014; Pan et al., 2012).

3.4.3 Virtual Character Ratings and Self-Other Overlap

Participants rated their feelings of rapport, trust and similarity towards each virtual character, as well as the smoothness of the interaction. They also rated feelings of self-other overlap towards the specific virtual character, virtual characters in general, their best friend and other people in general.

For each dependent variable, we conducted a two-way mixed-design ANOVA to test the within-participants effect of mimicry condition (mimicry vs. non-mimicry) and the between-participants effect of time delay (3s vs. 1s). The results are reported in Table 3-2. We found a significant main effect of mimicry on rapport in the expected direction (*p*

=.02). However, there were no other significant effects of mimicry or time delay, and no significant interactions.

To further explore the relationship between rapport and mimicry time delay, we calculated the effects of mimicry on rapport for the 1s and the 3s groups separately. We found that there was a significant effect of rapport in the 3s condition (mimicry M = .71, SD = .13; non-mimicry M = .64, SD = .13; $t(24) = 2.15$, p =.04, d =.49) but not in the 1s condition (mimicry M = .68, SD = .17; non-mimicry M = .62, SD =.16; $t(24) = 1.39$, p =.18, d =.36).

3.4.4 Bayesian Analyses

In order to further examine the evidence for mimicry effects, we carried out a series of Bayesian ANOVAs using JASP software, version 0.7.5.6 (JASP Team, 2016). Each Bayes factor (BF_{01}) indicates how much more likely the data were to occur under a null model, compared to the model described. The results (Table 3-3) show that participants' rapport ratings were more likely to occur under a model where mimicry has a main effect on rapport, compared to a null model. This could be considered weak or anecdotal evidence for the effect of mimicry on rapport (Jarosz & Wiley, 2014). There was ambiguous evidence for the main effect of mimicry on self-other overlap towards other people in general. There was weak evidence in favour of the null hypothesis for the main effects of mimicry on trust and similarity ratings. For all other main effects of mimicry and interactions between mimicry and time delay, the Bayes factors indicate substantial to strong evidence in favour of the null hypothesis.

Table 3-2. Effects of mimicry and time delay on dependent variables.

Measure	Mimicry *M* (*SD*)	Non-mimicry *M* (*SD*)	Main effect of mimicry			Main effect of time			Mimicry x time interaction		
			F(1, 48)	*p*	η_p^2	*F*(1, 48)	*p*	η_p^2	*F*(1, 48)	*p*	η_p^2
Rapport	.69 (.15)	.63 (.15)	5.54	.02	.10	.63	.43	.01	.002	.97	>.001
Trust	.65 (.14)	.62 (.15)	1.90	.17	.04	1.65	.21	.03	.31	.58	.01
Similarity	.55 (.17)	.51 (.20)	1.52	.22	.03	1.22	.28	.03	.04	.85	.001
Smoothness	.67 (.20)	.66 (.20)	.006	.94	>.001	.63	.43	.01	.03	.86	.001
Overlap: specific virtual character	.42 (.21)	.41(.19)	.28	.60	.01	3.60	.06	.07	.31	.58	.01
Overlap: virtual characters in general	.39 (.23)	.40 (.20)	.50	.50	.01	2.50	.12	.05	.12	.73	.002
Overlap: best friend	.60 (.26)	.60 (.26)	.21	.65	.004	.005	.94	>.001	.09	.77	.002
Overlap: others in general	.40 (23)	.42 (.23)	3.80	.06	.07	2.20	.14	.04	.05	.82	.001

Table 3-3. Bayes factors for the effects of mimicry and time delay on ratings.

Measure	Model BF_{01}			
	Mimicry main effect	Time main effect	Mimicry main effect + time main effect	Mimicry main effect + time main effect + mimicry x time interaction
Rapport	0.34	2.94	1.01	3.76
Trust	2.01	1.45	3.01	9.01
Similarity	2.34	1.99	4.69	16.85
Smoothness	4.79	2.02	9.85	34.74
Overlap: specific virtual character	4.26	0.72	3.21	9.14
Overlap: virtual characters in general	3.83	1.02	3.79	14.16
Overlap: best friend	4.41	1.44	6.66	22.47
Overlap: others in general	0.89	1.66	1.52	5.21

3.5 Discussion

In this study, we developed a new virtual mimicry paradigm and investigated whether the timing of mimicry affects how people respond to being mimicked. We programmed virtual characters to mimic participants' head and torso movements at a delay of three seconds or one second, or to move according to pre-recorded motion in a non-mimicry condition. We measured how many participants consciously noticed being mimicked, and the remaining participants rated their feelings of smoothness, similarity, rapport and trust towards the virtual characters. These data allow us to consider if being mimicked by a virtual character can act as social glue, and if the timing of the mimicry matters.

3.5.1 Can Virtual Mimicry act as Social Glue?

Our data give a reasonably positive response to the question of whether being mimicked leads to a more positive social evaluation of the mimicker. Participants rated a level of co-presence (feeling the virtual characters were really present with them) comparable to virtual reality experiments using more immersive technology (Friedman et al., 2014; Pan et al., 2012), suggesting our system achieved a strong degree of realism. Participants also gave higher rapport ratings to characters who mimicked compared to those who did not. This is consistent with previous research using virtual and human mimickers (Bailenson & Yee, 2005; Chartrand & Bargh, 1999; Stel et al., 2011) and supports the idea that mimicry can act as social glue even in tightly controlled VR settings. However, there are some caveats to this finding. First, we did not find effects of mimicry on any of our other dependent measures (trust, similarity, interaction smoothness and self-other overlap). It could be that the outcomes we chose to measure are more fragile than we expected; for example, other studies have also failed to find statistically significant effects of imitation on trust (Bailenson et al., 2008; Verberne et al., 2013), similarity (Reddish et al., 2013) and self-other overlap (Hogeveen et al., 2014). Second, the effect on rapport which we did find was small (η_p^2 = .10) and would not meet a Bonferroni correction. Bayesian analyses showed only weak evidence in favour of this effect existing in our data (BF_{10} = 2.94).

It could be argued that these findings reflect a failure of virtual mimicry to achieve the same effect sizes as naturalistic mimicry. We acknowledge that current virtual reality methods are not perfect, and participants in this study were aware they were interacting with a virtual agent and not another real person. However, our co-presence scores of 4.45 (Study 1) and 4.75 (Study 2) on a scale from 1-7 are comparable to other VR studies. Such studies have successfully replicated psychology effects in the domain of body ownership (Maister, Slater, Sanchez-Vives, & Tsakiris, 2015), and interactions with virtual characters have been able to successfully generate joint attention (Schilbach et al., 2009), proximity effects (McCall & Singer, 2015) and audience effects (Slater et al., 1999; Slater, Pertaub, Barker, & Clark, 2006) among other phenomena. This would suggest that virtual reality can generate socially realistic scenarios which replicate real-world psychological phenomena.

3.5.2 What is the Role of Timing in Mimicry?

The data gives conflicting accounts of the role of timing in mimicry. On the one hand, we found a significantly higher rate of mimicry detection at a delay of 1s compared to 3s, consistent with Bailenson et al. (2004). This implies that some cognitive process is able to detect when someone else mimics me, and this process is tuned to the timing of their movements. However, the present study cannot distinguish if this process is specific to mimicry, or simply sensitive to any contingent behaviour (Heyes, 2011). On the other hand, after people who consciously detected mimicry were excluded from the analysis, we found no significant differences in ratings after being mimicked with a 3 second delay compared to a 1 second delay (although the effect of mimicry on rapport only reached significance in the 3s group). This was surprising, and implies that any cognitive mechanism which responds positively to being mimicked outside awareness is not tuned to a narrow time window. Thus, the process of consciously detecting mimicry may be independent of how mimicry is used to evaluate another person. Though counterintuitive, this suggestion is consistent with results from a previous study which found the spatial correspondence between mimicker and mimickee movements affected mimicry detection but not the social evaluation of the mimicker (Bailenson et al., 2008). Therefore, it is possible that any mechanism which triggers a positive response to being mimicked has a broadly tuned time window for registering any mimicry. This possibility would fit with the observation that 'positive' mimicry effects could actually be driven by negative reactions to an absence of mimicry (van Baaren, Decety, et al., 2009), as well as evidence that people expect certain levels of contingency when interacting with another person (or robot) and respond positively as long as that threshold is met (Bigelow, 1998, 2001, Yamaoka, Kanda, Ishiguro, & Hagita, 2006, 2007).

3.6 Conclusion

In this study we developed a new virtual mimicry paradigm which achieved a good level of social realism and can be directly compared to traditional studies in which confederates covertly mimic participants or not during a social interaction task. In line with the social glue theory of mimicry, we found that participants rated significantly more rapport towards a virtual character who mirrored their head and torso movements, compared to a virtual character which moved in a pre-recorded way. However, there were no significant effects of mimicry on any of the other social outcomes we measured. These findings suggest that virtual mimicry is able to generate positive social effects, but that the effects of mimicry alone may be smaller than expected when we strictly control all other social signals such as smiling and eye contact. Varying the precise timing of mimicry between one and three seconds did not affect any social outcomes, although there was a higher rate of conscious detection at the one second delay. This suggests that the detection of mimicry may be tuned to the timing of the mimicker's actions, but this process may be independent of how positively we evaluate people who mimic us.

Chapter 4. Does mimicry hold as social glue across group boundaries? A test using virtual reality

4.1 Abstract

Virtual reality may provide a highly controlled method for strictly testing the 'social glue' theory which claims that mimicry leads to positive social outcomes such as rapport and trust. In this study, we aimed to replicate our previous finding that virtual mimicry leads to rapport, and to extend this by testing whether positive social effects of mimicry are modulated by the ingroup or outgroup status of the mimicker. To generate ingroup and outgroup mimicry, Forty participants from European or East Asian backgrounds interacted with four avatars, two of European appearance and two of East Asian appearance. Two avatars mimicked while the other two did not. During the interaction, participants rated how much they liked each character, and afterwards they rated feelings of rapport and trust. Participants also completed an immersive virtual maze task to measure implicit trust behaviour. Across all measures we found no main effects of mimicry and no interactions with group membership. These null results were calculated in line with a pre-registration. We conclude that being mimicked does not always increase rapport or trust when all other social signals are controlled.

4.2 Introduction

In Chapter 3, we established that our VR mimicry paradigm could generate increases in rapport, similar to naturalistic paradigms with confederate mimickers. The present study aimed to replicate this initial finding with a rigorous, pre-registered procedure (Hale & Hamilton, 2015) and to further test if the social consequences of mimicry may be altered by group membership. We consider this study to be the more definitive test of the 'social glue' hypothesis that being mimicked leads to liking and positive evaluations of the mimicker. In this study we also measure implicit trust behaviour towards mimicking and non-mimicking virtual characters using the virtual maze task introduced in Chapter 2.

While it has been suggested that mimicry is an 'honest signal' used in all human societies (Pentland, 2010), it is not clear whether mimicry leads to increases in rapport and trust between people from different cultural groups (Chartrand & Lakin, 2013). This is an important question, because many proposed applications of mimicry may involve cross-cultural settings or other ingroup-outgroup relationships which could alter the expected positive effects of mimicking. For example, in a teaching context (Bernieri, 1988; LaFrance & Broadbent, 1976), mimicry might not have the anticipated benefits for rapport if the teacher has a different cultural background to their students. When the mimicker and mimickee come from different backgrounds this could also change how mimicry affects rapport, persuasion and disclosure in settings such as business sales and negotiations (Maddux et al., 2008; Tanner, 2008; van Swol, 2003) and investigative interviews (Abbe & Brandon, 2013a, 2013b; Stel et al., 2009).

There have been very few previous studies addressing whether mimicry can be used to build rapport and trust between members of different social groups. Studies focusing on the *mimicker* show that people typically produce less mimicry towards others who they initially dislike (Stel, Blascovich, et al., 2010), competitors (LaFrance, 1985), outgroup members (Bourgeois & Hess, 2008; Yabar et al., 2006) and others from a different race (Johnston, 2002). Similar effects have also been found for the production of facial mimicry towards outgroup members (Bourgeois & Hess, 2008; Mondillon, Niedenthal, Gil, & Droit-Volet, 2007; van der Schalk et al., 2011) as well as behavioural synchrony (Miles, Lumsden, Richardson, & Macrae, 2011). On the other hand, a small number of mimicker studies have found that mimicking an outgroup member can lead to more liking and empathy towards the outgroup (Inzlicht, Gutsell, & Legault, 2012; Szuster & Wojnarowska, 2016; van der Schalk et al., 2011). However, there is very little research about how *mimickees* evaluate mimicry from outgroup members. Some studies suggest being mimicked by an outgroup member depletes cognitive resources, making it harder to perform a Stroop task (Dalton et al., 2010) and leads to feeling colder when estimating the temperature of a room (Leander et al., 2012). On the other hand, a recent experiment found that being mimicked by an outgroup member (a Palestinian virtual character) led participants (Jewish Israelis) to increase empathy towards their outgroup in terms of the language they used (Hasler et al., 2014). Finally, a study of romantic partners found that if one partner was mimicked by a stranger, this could lead to stronger bonding between the two partners (Kouzakova, Karremans, et al., 2010). These mixed findings suggest that being mimicked by an outgroup member can both challenge people's expectations and also have some positive social outcomes.

Virtual reality has several advantages for manipulating the ingroup or outgroup status of a mimicker, especially on the basis of culture. Firstly, the same computer algorithm can be used to generate exactly the same mimicry (and non-mimicry) behaviour towards all participants. This avoids the possibility that human confederates, who cannot be blind to the participant's cultural group, might subtly and unconsciously change their behaviour towards different groups of participants. For example, people tend to align their accent and lexical patterns with people from a different culture (Bock, 1986; Giles & Powesland, 1975; Niederhoffer & Pennebaker, 2002), but virtual characters are not affected by this tendency. Secondly, virtual reality makes it very straightforward to carry out repeated-measures research designs where the cultural group of the mimicker is counterbalanced simply by switching the computer code, without the need to train multiple confederates as mimickers and non-mimickers. Thirdly, virtual reality has been useful for studying sensitive ingroup-outgroup scenarios, which might otherwise be hard to study with human confederates for ethical reasons. For example, Slater et al. (Slater et al., 2013) looked at how participants responded to a violent altercation between rival football team supporters, and Hasler et al. (2014) studied how mimicry could affect empathy between Israelis and Palestinians in the context of a politically-charged debate.

The present study tested whether mimicry leads to rapport and trust across individuals from the same social group (in-group pair) and different groups (out-group pair). To create in-group and out-group pairs, we used virtual characters with different apparent ethnicities and recruited participants from East Asia and Europe. Each participant met with 4 virtual characters who mimicked or did not and who did or did not appear to share the participant's regional background, in a 2x2 factorial within subjects design. Mimicry was implemented with a 3 second time delay, and we continued testing until we had complete datasets from 40 participants without any detection of mimicry. This sample size was chosen following a power analysis based on our previous study, in which we found a small effect of mimicry ($\eta_p^2 = .10$). We calculated that a sample size of 40 could detect the same size effect with power of .99, or a smaller effect size ($\eta_p^2 = .06$) with power of .95.

Participants completed a photo description task with each virtual character, during which they were exposed to mimicry or non-mimicry behaviour. They rated likeability, trust and rapport towards each virtual character, and we also measured behavioural trust using the virtual maze task introduced in Chapter 2 (note the maze task was developed after completing our previous mimicry study in Chapter 3). In this study, we did not include measures of self-other overlap, interaction smoothness or feelings of similarity, because we found null effects of mimicry on these outcomes in our previous study (Chapter 3). We predicted that participants should show a main effect of mimicry, with greater liking, trust and rapport towards virtual characters which mimicked. Furthermore, we predicted that this mimicry effect should be modulated by cultural group, such that all effects would be more strongly positive for in-group virtual characters compared to out-group virtual characters. I report the results as specified in our pre-registration (Hale & Hamilton, 2015), and clearly label additional exploratory analyses which were not specified in the pre-registration.

4.3 Method

4.3.1 Participants

Participants completed a screening questionnaire to ensure they were suitable for the study (Table 1). We only selected participants if their answers to the screening questionnaire indicated they self-identified strongly with one of the groups chosen for the study. We included (1) individuals who identified strongly as European and (2) individuals who identified strongly as East Asian, had spent less than 1 year in the UK, and had spent most of the last 10 years in an East Asian country. The extra criteria for the East Asian group aimed to minimise overlap caused by recruiting and conducting the experiment in a European country.

Fifty-four participants (37 female, M_{age} = 25.14 years, *SD* = 3.66 years) out of 236 total respondents were selected to take part in the study. Ethical approval was obtained from the UCL Institute of Cognitive Neuroscience Ethics Chair (Project ID Number ICN-AH-PWB-3-3-14a) and the Ministry of Defence Research Ethics Committee (564/MODREC/14). The methods were carried out in accordance with the approved guidelines and the Declaration of Helsinki (2013). All participants gave written informed consent to take part.

4.3.2 Virtual Mimicry

We used the same equipment and algorithm as in the previous study (Chapter 3) to generate virtual mimicry. We prepared five virtual characters for the study (Figure 4-1). We used the same practice virtual character (a male named Mike) from the previous study (Chapter 3). We also used the same Anna and Becky virtual characters as the previous study, because they had European appearances and their voices were recorded by volunteers with native British accents. The other two virtual characters, named named Su Lin and Tian Tian, had an East Asian appearance and their voices were recorded by volunteers speaking English with native Chinese accents. All virtual characters were programmed to blink, make eye contact with the participant and speak

Table 4-1. Screening questionnaire and selection criteria.

Question	Response	Selection criterion for European Group	Selection criterion for East Asian Group
What is your nationality?	Text field	European nationality	East Asian nationality
How much is your nationality important to your identity?	1-7 Likert scale: *extremely unimportant* – *extremely important*	Score >= 6	Score >= 6
What is your ethnicity?	Text field		
How much is your ethnicity important to your identity?	1-7 Likert scale: *extremely unimportant* – *extremely important*		
What is your first language(s)?	Text field	Must have English as a first language **OR** speak fluent English	Must have English as a first language **OR** speak fluent English
How much is your first language(s) important to your identity?	1-7 Likert scale: *extremely unimportant* – *extremely important*		
What other languages do you speak? Please list any languages you speak well enough to hold a conversation.	Text field	Must have English as a first language **OR** speak fluent English	Must have English as a first language **OR** speak fluent English
How many years have you lived in the UK? Please include time in the past as well as your current residence.	Less than 1 year 1 - 5 years 5 - 10 years 10 years or more		Must have lived in the UK for less than 1 year.
Over the last 10 years, which country have you lived in for the most time?	Text field		Must have lived in an East Asian country for most of the last 10 years
Date of Birth	DD/MM/YY	Aged 18-35	Aged 18-35
Gender	Male / Female		

according to pre-recorded sound files, as in the previous study.

4.3.3 Experiment Procedure

Each participant sat at the desk in front of the projector screen, and the experimenter fitted the participant with motion tracking sensors. First the participant practiced the photo description task with the practice character, who never mimicked them.

Following the practice, the participant completed the four experimental conditions in a counterbalanced order. The conditions were: (1) mimicry from an in-group member, (2) non-mimicry from an in-group member, (3) mimicry from an out-group member and (4) non-mimicry from an outgroup member. Half the participants met virtual characters from their in-group first, and the other half met virtual characters from their out-group first. For half the participants, the first character in the group mimicked them, and for the other half the second character in the group mimicked. The order of characters in each group was always Anna then Becky, and Su Lin then Tian Tian. An example order of conditions is given in Figure 4-2, which summarises the experiment procedure.

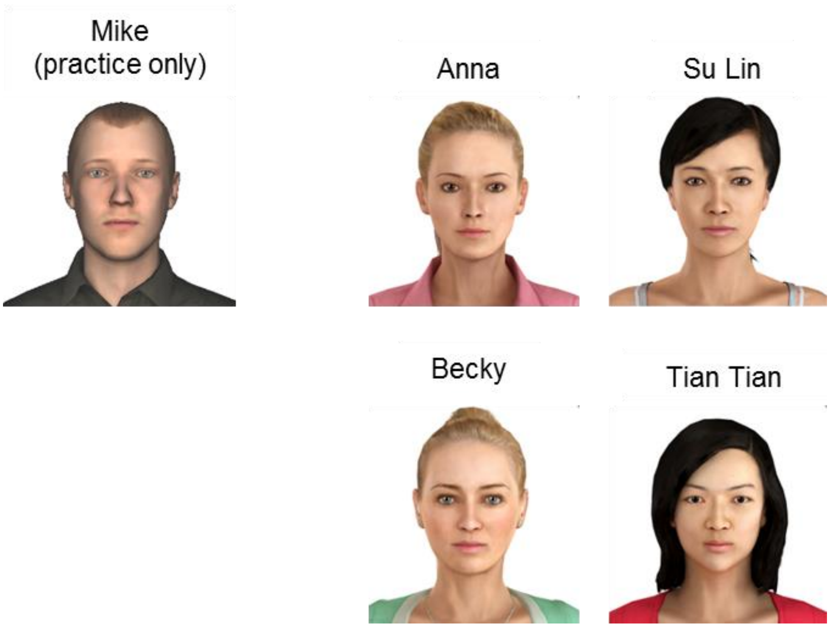

Figure 4-1. Virtual character appearances.

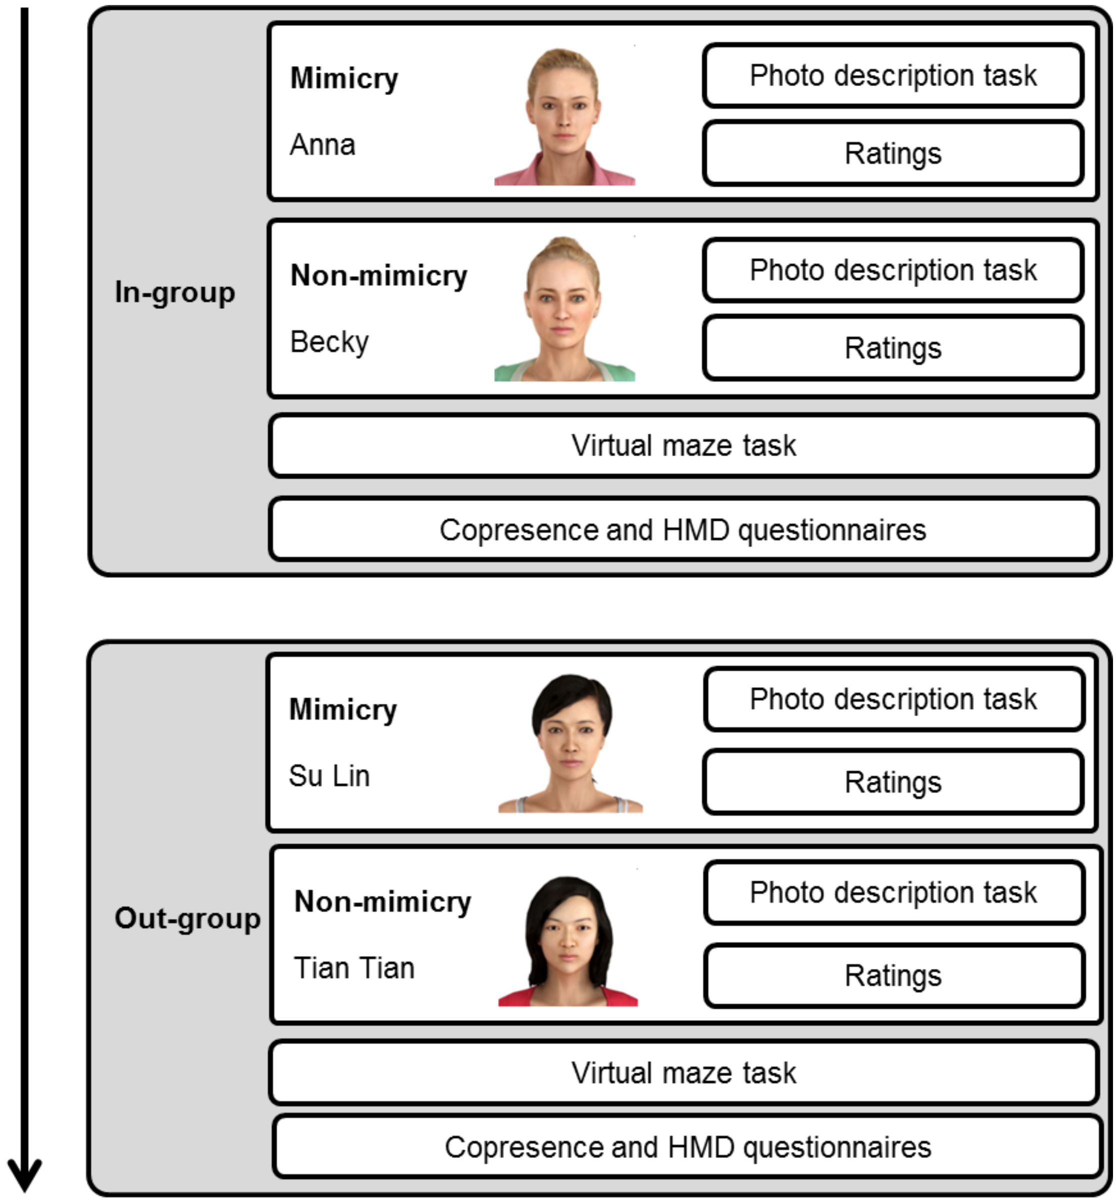

Figure 4-2. Experiment procedure, showing an example order of conditions. Group and mimicry were counterbalanced across participants.

Photo description task (mimicry manipulation). First, the participant completed the mimicry induction (photo description task) with the first virtual character (e.g. Anna), as in the previous study (Chapter 3). During the photo description task, the virtual character either mimicked the participant's movements after a 3s delay, or performed pre-recorded movements (as described in Chapter 3). Half-way through the photo description task, the participant was prompted to rate how much they agree that with the statement *'I think Anna is very likeable'* on a continuous scale from strongly disagree to strongly agree. We introduced this additional rating in the present study in case an effect of mimicry could be detected *during* the interaction which could not be detected afterwards, in light of the weak effects of mimicry on post-interaction ratings found in Chapter 3.

Virtual character ratings. Next, the participant rated their feelings of rapport and trust on a sliding scale from *'strongly disagree'* (scored 0) to *'strongly agree'* (scored 1). The rapport and trust items were the same as those in the previous study.

Then the participant repeated the photo description task and ratings with the second virtual character (e.g. Becky). The first two characters were always a mimicker and non-mimicker from the same group.

Virtual maze task. After the participant completed the photo task and ratings with the first two characters, the experimenter removed the motion sensors and prepared the participant for the virtual maze task. Then the participant completed the virtual maze task with the two virtual characters they had interacted with so far (e.g. Anna and Becky). In the task, the participant navigated through a series of rooms and in each room they could ask for advice from Anna and/or Becky about which way to go next. The maze was displayed via an Oculus Rift DK2 head-mounted display (HMD). This device allows the participant to look around a virtual 3D space as if they are really there. The participant was provided with a joystick to move through the virtual space, similar to playing a video game. Participants first completed some practice to familiarise them with the 3D environment. Then they were instructed to find the way out of the maze through the least number of rooms possible. The computer program ensured that each participant continued until they had gone through twelve rooms (corresponding to twelve trials). For full details of the trial stimuli and procedure, see Chapter 2. On each trial we recorded whether the participant approached each character for advice or not, and whether they followed the advice of each character or not. This provided two measures of trust: (1) how often each character was approached (expressed as a percentage of trials) and (2) how often each character's advice was followed (expressed as a percentage of trials).

Co-presence and VR questionnaires. At the end of the virtual maze task participants completed a questionnaire about their experience of co-presence, as in the previous study. Then they were asked to complete a short questionnaire about their experience in the virtual maze and whether they experienced any symptoms of motion sickness, headache or eye strain. Data from the VR questionnaire are not included in the present study. Participants who reported any of the symptoms were not required to continue with the study.

Following the VR questionnaire, the participant repeated the entire procedure with the remaining two virtual characters (e.g. Su Lin and Tian Tian).

Screening questionnaire and debrief. At the end of the experiment the participant completed the same screening questionnaire that we used at the recruitment stage (Table 4-1). This was to make sure that participants still reported background characteristics that met our group selection criteria. Finally, the participant also provided written and verbal answers to a series of questions to determine whether they had guessed the purpose of the experiment or noticed the mimicry manipulation.

4.3.4 Contributions

Joanna Hale implemented the virtual mimicry algorithm in Vizard (this does not include inverse kinematic modelling which was done automatically in MotionBuilder), recorded all audio stimuli and programmed the photo description task and virtual maze task. Joanna Hale carried out all participant recruitment, data collection and analyses.

4.4 Results

Six participants were excluded from analyses due to technical failure of the motion tracker. Eight out of the remaining 48 participants (16.6%) detected the virtual mimicry manipulation. All of the detectors were excluded from the analyses, so the remaining sample was N = 40. Four participants did not complete the virtual maze task with every virtual character due to feelings of motion sickness, so their data was not included in analyses on that task.

4.4.1 Co-Presence

Participants completed co-presence ratings once after interacting with ingroup virtual characters (*M* = 4.58, *SD* = 1.21) and once after interacting with outgroup virtual characters *M* = 4.63, *SD* = 1.29). There was no significant difference in feelings of co-presence towards ingroup and outgroup members ($t(39) = .43$, $p = .67$, $d = 0.09$). Therefore we averaged the two scores together. The median co-presence score was 4.75 (*M* = 4.61, *SD* = 1.19, *Range* = [1.63, 6.75]), consistent with our initial study, reported in Chapter 3, and other VR research (Friedman et al., 2014; Pan et al., 2012).

4.4.2 Liking, Rapport and Trust Ratings

We conducted two-way, repeated-measures ANOVAs to test the effects of mimicry and group membership on ratings of liking, rapport and trust towards each virtual character. We did not find any significant main effects of mimicry or group membership, or any significant interaction effects (Table 4-2). There was a marginal effect of mimicry on trust in the expected direction: participants rated slightly more trust towards mimickers than non-mimickers.

4.4.3 Virtual Maze Task

We conducted two-way, repeated-measures ANOVAs to test the effects of mimicry and group membership on how often participants approached a virtual character in the virtual maze, and how often they followed a virtual character's advice (each expressed as a percentage of trials). We did not find any significant main effects of mimicry or group membership, or any significant interaction effects (Table 4-2). There was a marginal interaction between mimicry and group membership on how often each virtual character was approached in the maze: participants approached the mimicking ingroup member on a greater percentage of trials than the non-mimicking ingroup member, whereas this trend was reversed for outgroup members.

Table 4-2. Effects of mimicry and group membership on ratings and virtual maze task.

	Ingroup		Outgroup		Main effect of mimicry			Main effect of group			Mimicry x group interaction		
Measure	**Mimicry *M* (*SD*)**	**Non-mimicry *M* (*SD*)**	**Mimicry *M* (*SD*)**	**Non-mimicry *M* (*SD*)**	***F*(1, 48)**	***p***	**η_p^2**	***F*(1, 48)**	***p***	**η_p^2**	***F*(1, 48)**	***p***	**η_p^2**
Liking rating	.69 (.21)	.65 (.21)	.69 (.18)	.66 (.17)	2.35	.13	.06	.06	.81	.002	.13	.72	.003
Rapport rating	.72 (.16)	.70 (.15)	.70 (.15)	.69 (.13)	1.06	.31	.03	1.35	.25	.03	.25	.62	.006
Trust rating	.68 (.15)	.66 (.16)	.70 (.14)	.65 (.13)	3.81	.06	.09	.04	.84	.001	.43	.52	.01
Approach (maze)	.86 (.22)	.80 (.22)	.80 (.25)	.84 (.22)	.19	.67	.005	.05	.82	.001	3.08	.09	.08
Follow advice (maze)	.64 (.22)	.61 (.20)	.62 (.25)	.64 (.18)	.04	.85	.001	.06	.81	.002	.37	.55	.01

4.4.4 Exploratory Analyses

In addition to the pre-registered analyses described above, we conducted further exploratory analysis of this dataset.

Firstly, we carried out Bayesian ANOVAs using JASP software, version 0.7.5.6 (JASP Team, 2016) to test the main effects of mimicry, group membership and the interaction between mimicry and group membership on each of our dependent variables (Table 4-3). The results showed that our data favour a null main effect of mimicry on ratings of liking, rapport and trust, as well as trust behaviour in the virtual maze (all $BF_{01} > 1.6$). The marginal interaction between mimicry and group membership on how often participant's approached the virtual characters in the maze (reported above) had a Bayes factor of 5.37 in favour of the null hypothesis. Bayes factors favoured the null hypothesis for all main effects of group membership (all $BF_{01} > 3.5$) and all interactions between mimicry and group membership (all $BF_{01} > 30$).

Table 4-3. Bayes factors for the effects of mimicry and group membership on ratings and virtual maze task.

	Model BF_{01}			
Measure	**Mimicry main effect**	**Group main effect**	**Mimicry main effect + group main effect**	**Mimicry main effect + group main effect + group x time interaction**
Liking rating	2.11	5.79	11.70	49.67
Rapport rating	3.74	3.54	13.26	47.61
Trust rating	1.64	5.61	9.67	30.61
Approach (maze)	5.37	5.45	29.01	34.93
Follow advice (maze)	5.60	5.65	32.01	94.82

Note. All models have a prior model probability of 0.2 and include subject.

Secondly, in order to check whether participants' experience of co-presence modulated any of the effects in the pre-registered analyses, we carried out a series of ANCOVAs with co-presence as a covariate. When controlling for co-presence, there were no significant main effects of mimicry on any of our dependent variables (Table 4-4). There were also no significant main effects of group membership or any significant interactions between mimicry and group membership.

Thirdly, it is possible that we did not find a significant effect of mimicry due to fatigue effects: participants might have become bored or disengaged by the time they interacted with the second group of virtual characters. In order to investigate this possibility, we carried out a series of ANOVAs which only included data from the first group of virtual characters, i.e. the first half of the experiment. Mimicry was a repeated-measures factor and group membership was a between-subjects factor. There were no significant main effects of mimicry or group membership and no significant interactions between mimicry and group membership in this data (Table 5). This suggests there was not a mimicry effect from the first two trials which became masked by boredom or fatigue in later trials. Note, however, that this analysis is underpowered for testing between-subjects effects of group membership.

Table 4-4. Effects of mimicry and group on ratings and virtual maze task, controlling for levels of co-presence.

	Main effect of mimicry			Main effect of group			Mimicry x group interaction		
Measure	$F(1, 48)$	p	η_p^2	$F(1, 48)$	p	η_p^2	$F(1, 48)$	p	η_p^2
Liking rating	0.05	.83	.002	.14	.71	.004	3.03	.09	.09
Rapport rating	1.14	.29	.03	.12	.73	.004	.35	.55	.01
Trust rating	1.65	.21	.05	.88	.34	.03	.44	.51	.01
Approach (maze)	0.17	.68	.005	.12	.73	.004	3.86	.06	.11
Follow advice (maze)	0.006	.94	>.001	.008	.93	>.001	.58	.45	.02

Table 4-5. Effects of mimicry and group in the first half of the study.

Measure	Main effect of mimicry			Main effect of group			Mimicry x group interaction		
	***F*(1, 48)**	***p***	**η_p^2**	***F*(1, 48)**	***p***	**η_p^2**	***F*(1, 48)**	***p***	**η_p^2**
Liking rating	.17	.68	.004	0.13	.98	>.001	2.76	.11	.07
Rapport rating	.47	.50	0.01	.016	.90	>.001	.29	.60	.008
Trust rating	3.22	.08	0.08	.013	.91	>.001	.32	.57	.008
Approach (maze)	.01	.93	>.001	2.835	.10	.07	.49	.49	.01
Follow advice (maze)	.03	.87	.001	2.711	.11	.07	.02	.89	.001

Finally, because we found a marginal interaction between mimicry and group membership on how often participants approached each virtual character in the virtual maze, we further explored participants' approach behaviour using a less conservative measure: how often each virtual character was approached *first*. This is a less conservative measure of participants' trust towards each virtual character because it reflects a binary choice (e.g. Anna *or* Becky), and thus discriminates between different virtual characters more than our original measure of how often each virtual character was approached *at all*, which could be both virtual characters (e.g. Anna *and* Becky) on a given trial. A traditional ANOVA on this less conservative approach measure, with mimicry and group membership as factors, revealed a significant interaction between mimicry and group membership in the same direction as the marginal interaction for the more conservative measure (F (1,35) = 4.64, p = .04, η_p^2 = 0.12). The ingroup mimicker (M = .53, SD = .22) was approached first more often than the ingroup non-mimicker (M =.43, SD = .23), and the opposite was true for outgroup members (mimicker M = .46, SD = .27; non-mimicker M = .51, SD = .27). However, Bayesian analysis for this effect strongly favoured the null hypothesis (BF_{01} = 27.2) and also favoured the null hypothesis for the main effect of mimicry (BF_{01} = 5.0) and the main effect of group (BF_{01} = 5.6).

4.5 Discussion

In this study participants were mimicked (or not) by a virtual character from their ingroup or an outgroup during a photo description task which roughly simulated a conversation. We measured the effect of mimicry on ratings of liking, rapport and trust, as well as implicit trust behaviour in a virtual maze task, and pre-registered our analyses. With this rigorous test of the effects of virtual mimicry, we found null results. In the pre-registered analyses we did not find any significant main effects of mimicry or group membership, or any significant interactions, although there were some trends in the predicted directions. In particular, we found a marginal interaction between mimicry and group membership on how often each virtual character was approached in the virtual maze. An additional exploratory analysis found the interaction was significant for a less conservative measure, being approached *first* (rather than at all). Although this might tentatively be taken as evidence that people respond negatively (instead of positively) to mimicry from an outgroup member, it is important to note this exploratory finding would not meet a Bonferroni correction for six comparisons. Furthermore, Bayesian analysis showed that our data provides evidence in favour of the null hypothesis for all main effects of mimicry and group membership, and all interactions. Therefore, we are not able to draw any strong conclusions about how cultural group membership may modulate mimicry effects. Although it would be tempting to speculate that mimicry effects are reversed for outgroup mimickers based on some of the trends in our data, we emphasise that there is very little evidence for this modulation and that much more conclusive research would be needed to make any recommendations about using mimicry to build rapport or trust in cross-cultural contexts.

These null results contrast with our initial exploratory study (Chapter 3), which found that participants rated significantly higher feelings of rapport towards a virtual character that mimicked them with either a 1 second or 3 second delay, compared to a virtual character with pre-recorded movement. Inspection of the data divided by time delay showed that mimicry increased rapport at 3 seconds delay but not at 1 second.

This is noteworthy because it means that the present study did not use the wrong time delay to find positive effects of mimicry. Additional analyses also suggest our null results were not due to fatigue effects. Whereas our initial exploratory study had many data analysis options available, the present study followed a single analysis pathway which was preregistered, and therefore we consider the present study to be more definitive in providing a clear test of the hypothesis that mimicry leads to liking, rapport and trust. This suggests that the result from our initial study may have been a false positive, and, taking both studies together, our data do not provide strong evidence that being mimicked leads to positive social evaluations of the mimicker.

Our data is in line with previous studies of mimicry using VR, which have reported mixed results. Bailenson and Yee's (2005) virtual mimicry system had a positive effect on participants' impression of a virtual character, measured using a scale that taps into congeniality, knowledgeability and sincerity. Vrijsen et al (Vrijsen, Lange, Dotsch, et al., 2010) used the same paradigm in a study comparing socially anxious and non- anxious women's responses. They did not report the main effect of mimicry, but report a near-significant positive effect of mimicry on how non-anxious women evaluated the virtual character in terms of likeability and friendliness. Verberne et al (2013) used the same mimicry algorithm and found inconsistent results across a range of measures. With one virtual character, mimicry had a significant positive effect on liking and trust, but with another virtual character the same mimicry manipulation did lead to significant positive effects. In a subsequent study in which the mimicking virtual character also had similar goals and appearance, Verberne et al. (2015) also found inconsistent effects across different measures of trust. It is worth noting that ratings of liking, trust and self-other overlap also show inconsistent effects of mimicry in traditional research settings where human confederates were trained to mimic participants (Hale & Hamilton, 2016a). Therefore the positive social effects of being mimicked may be more subtle or fragile than is generally assumed. Overall, we conclude that the mimicry of participant head and torso movements alone by a virtual character does not lead to substantial increases in rapport and trust.

4.5.1 Implications for the Social Glue Theory of Mimicry

The social glue theory of mimicry (Dijksterhuis, 2005; Lakin et al., 2003) suggests that non-conscious mimicry occurs across a wide range of behaviours including speech patterns, facial expressions, emotions, postures, gestures and mannerisms, and that mimicry of any of these behaviours can lead to positive social consequences such as rapport and interpersonal closeness. Virtual reality provides a strong test of this hypothesis, because it allows us to specifically control one type of mimicry while keeping all other social behaviours (speech, gaze, facial expression) exactly the same. Our findings across the present study and the previous study (Chapter 3) thus cast doubt over a strong version of the social glue theory, in which all types of mimicry must have positive social effects. It seems that in the specific case of mimicking head and torso movements in a mirror fashion, mimicry does not lead to increases in liking, rapport or trust.

It is possible that our virtual characters lack other critical behaviours and did not achieve a realistic replication of natural mimicry. Adding other social behaviours such as nodding and more facial expressions, might lead to stronger effects of virtual mimicry (Gratch, Wang, Gerten, Fast, & Duffy, 2007) by making it more similar to real life mimicry. This may involve a variety of other social signals such as emotional imitation (Chartrand et al., 2005; Wallbott, 1995), turn-taking (Pentland, 2010; Wallbott, 1995) and eye-contact (Wang, Newport, et al., 2011). Although traditional studies have attempted to control for the confounding effects of these signals by videotaping and rating confederates' behaviour (Maddux et al., 2008; Stel, Blascovich, et al., 2010; Stel et al., 2009), it is likely that some extra social cues still accompany mimicry in naturalistic paradigms. Therefore by controlling all other social signals, virtual mimicry may achieve smaller effect sizes than naturalistic mimicry (although in the exploratory study we found an effect on rapport with eta squared of .1, similar to naturalistic studies). It is also possible that mimicry may have different effects in the context of our structured conversation task, compared to other contexts involving different levels of interaction (e.g. a free conversation; Chartrand & Bargh, 1999; van Baaren et al., 2004) and different external cues (e.g. a negotiation agenda; Maddux et al., 2008). However, in order to test the claim that mimicry itself leads to positive social effects, virtual mimicry paradigms like ours provide the strictest test without other interfering social signals. Despite some positive results from VR scenarios which support the social glue theory, our data provides evidence in favour of null effects of virtual mimicry on rapport and trust.

Within the context of the wider reproducibility crisis in psychology and the mixed evidence for positive effects of mimicry reviewed in Chapter 1, our studies contribute novel evidence to debates around the social glue theory. There have been very few previous studies of virtual mimicry and those reported show mixed results (Bailenson & Yee, 2005; Hasler et al., 2014; Verberne et al., 2013a). We extended Bailenson and Yee's (2005) paradigm by implementing mimicry of body position as well as head motion, and by manipulating mimicry within the context of an interactive social exchange rather than a one-sided monologue. Our studies go demonstrate that the significant effects reported in previous virtual mimicry studies are difficult to replicate using novel methods and pre-specified hypothesis tests. Our Bayesian analyses also go beyond previously reported results to provide evidence in favour of null effects of mimicry. In addition, our studies have extended our wider understanding of the dowenstream effects of mimicry by capitalising on the advantages of virtual mimicry to show that the timing of mimicry and the group membership of the mimicker do not significantly modulate people's responses to being mimicked by a virtual character.

4.6 Conclusions

Across two tightly controlled studies in virtual reality (Chapter 3 and Chapter 4), we found that mimicry did not reliably lead to increases in rapport and trust when the delay was 3 seconds or 1 second, or when the mimicker came from an ingroup or outgroup. We cannot overturn a large body of work based on two studies, but our findings emphasise the fragile nature of mimicry effects, and suggest that further work with larger sample sizes and rigorous methods will be needed to determine if mimicry really is social glue. Looking forward, further exploration and quantification of natural mimicry behaviour during social interactions would be valuable for developing more realistic virtual characters that are able to achieve larger effects of mimicking.

In the next chapter, we move on to investigate mimicry from a fresh perspective using a different methodological approach. While virtual mimickers offer much tighter control than human confederates, it is currently difficult to achieve completely naturalistic mimicry from virtual characters. Therefore in Chapter 5 we explore a complementary approach and investigate the parameters of spontaneous coordination between dyads of naïve participants. Representing an initial step in a new direction, we focus on quantifying these parameters rather than testing the downstream consequences of spontaneous mimicry.

Chapter 5. Get on my wavelength: Interpersonal coordination in naturalistic conversations

5.1 Abstract

We currently have little data about the precise parameters of mimicry in natural conversations, and the role of mimicry in relation to other types of coordination remains unclear. Recently, wavelet analysis has been identified as a tool to measure interpersonal coordination in motion data recorded from naturalistic social interactions. In a pilot study, we used high-resolution motion capture to record head and torso movements of 20 dyads over multiple 90 second conversations. Following exploratory analyses on the pilot dataset, we replicated the same procedure with a new sample of 31 dyads and pre-registered our analyses of the new dataset. We used wavelet analysis to examine levels of mimicry across a spectrum of different motion frequencies. We tested levels of mimicry in real vs. pseudo interactions, created by shuffling the dataset. The results suggest that natural conversations are characterised by greater than chance coordination at frequencies below 1.5 Hz. Surprisingly, real interactions also showed *less* than chance coordination of head movements at frequencies of 1.5 – 5 Hz, suggesting that systematic *decoupling* of head movements occurs in natural conversations alongside spontaneous mimicry. Levels of mimicry did not differ between monologue and dialogue interactions or following prosocial vs. antisocial primes. We discuss the implications of our data and suggest future directions.

5.2 Introduction

In the previous two chapters, we found that precise mimicry of head and torso movements by a virtual character did not have any reliable effects on rapport or trust. This could be because virtual mimicry is currently unable to replicate spontaneous and reciprocal patterns of mimicry that may occur between two people in a natural interaction. In virtual mimicry the virtual character always exactly matches the participant like a mirror with a fixed delay, but natural mimicry may involve switching between mimicker and mimickee roles, as well as variation in the degree to which people match their movements in time and space. Natural mimicry also occurs in a rich context of other coordinated actions, nonverbal signals and speech. It has traditionally been challenging to measure and model mimicry in these contexts, but modern motion capture technologies and analysis techniques make it possible to measure natural mimicry in precise detail (Chapter 1). Therefore, in the present study we move away from controlled manipulations of mimicry and use these methods to explore the natural parameters of interpersonal coordination in face-to-face interactions. In the present chapter, we focused on quantifying these parameters rather than measuring downstream effects of mimicry on rapport, trust or any other social outcomes.

Interpersonal coordination is an umbrella term that covers many different kinds of coordinated action, including mimicry, synchrony, deliberate imitation and complementary movements. When the term ‘interpersonal coordination’ was first coined by Bernieri & Rosenthal (Bernieri & Rosenthal, 1991), they distinguished between behaviour matching (mimicry) and behavioural synchrony. Since then, mimicry and synchrony have occupied separate research literatures and discussions have focused on similarities and differences in their causes and effects (Chartrand & Lakin, 2013; Hove & Risen, 2009; Piercarlo Valdesolo et al., 2010). Other sub-categories of interpersonal coordination such as imitation, entrainment and contingent behaviour also tend to be studied in separate fields. For experimental purposes, the distinctions between these sub-categories are important for generating and testing specific hypotheses. However, for modelling realistic social interactions, treating them separately may be less helpful. Everyday conversations, meetings, games, sports, music, teaching and other interactions all involve a dynamic mixture of mimicry, synchronisation, complementary and joint actions which may not be easily disentangled. In addition, it is unclear whether sub-categories of interpersonal coordination form a single continuum dependent on the same basic cognitive mechanisms, or if there are clear distinctions in the mechanisms for different interpersonal behaviours (Hale & Hamilton, 2016a). Therefore, in order to understand the natural production of mimicry, we go beyond a strict definition of mimicry in this study and explore wavelet analysis as a method for modelling interpersonal coordination more broadly.

Natural social interactions also involve reciprocal patterns of coordinated behaviour. This is emphasised in synchrony research, which typically measures synchronisation at the level of dyads or groups (Lakens et al., 2016). In contrast, research on mimicry has typically involved experiments in which the participant is either mimicker or mimickee. This approach, where the individual is the unit of analysis, has generated a lot of research into what causes the mimicker to mimic, and (to a lesser extent) the downstream consequences of mimicry for the mimickee. Therefore, our current understanding of spontaneous mimicry comes from observing how often one person matches a target's movements within a critical time window (Stel et al., 2009; Stel et al., 2010). Very few studies have attempted to measure spontaneous mimicry as a reciprocal process within a dyad (or larger group) or test causes and consequences of mimicry at this level of interaction. However, this could be critical for determining naturalistic mimicry parameters such as its timing or rhythmic properties. In particular, research focusing on the organisation of nonverbal signals in social interactions suggests that natural patterns of social exchange within a participant dyad cannot neccesarily be captured when one participant is replaced with an experimenter or confederate (Chapter 1; Bavelas & Healing, 2013; Brown-Schmidt & Tanenhaus, 2008; Kuhlen & Brennan, 2013). Therefore, in the present study we recruited participant dyads and aimed to explore their patterns of interpersonal coordination at the dyadic level.

5.2.1 Recording Dyadic Social Interactions

Researchers have used several approaches to measure interpersonal coordination in dyads (Lakens et al., 2016). Early research on interpersonal coordination involved videotaping natural conversations, often in clinical contexts between therapists and their clients. The videotapes were then extensively coded, frame-by-frame in order to record changes in speech and posture (Bernieri, Gillis, Davis, & Grahe, 1996; Condon & Ogston, 1966; Kendon, 1970; LaFrance & Broadbent, 1976). Scheflen (1964) noted that two (or more) people interacting in therapy sessions often adopted congruent postures, either directly matching or mirroring one another’s body positions. Shortly after that, detailed video analyses of conversations by Condon & Ogston (1966) and Kendon (1970) found that people tended to move different or similar body parts in coordination both with each other and with their own speech rhythms at the level of words, syllables and phonemes. Central to this microanalysis was the idea that synchrony could occur between different behaviours, which required intensive coding of many movements from each person interacting at high time resolution (up to 48 frames per second). Although this provided very rich data, slow hand annotation was a major limitation (Grammer, Kruck, & Magnusson, 1998) and it is likely that chance coordination would occur when so many different behaviours are coded (Cappella, 1981; McDowall, 1978). In addition, even trained coders with high internal reliability might have introduced bias into the scoring of coordinated movements (Lumsden et al., 2012).

Subsequent research relied on untrained observers to rate the amount of coordination in videotaped social interactions. This approach relies on the gestalt assumption that people can accurately perceive stimulus properties such as the level of coordination between two people without the need for training (Bernieri & Rosenthal, 1991). In order to test the validity of this approach, Bernieri et al. (Bernieri, 1988; Bernieri et al., 1988) gave observers videotapes of real interactions and ‘pseudo interactions’ where interaction partners from different videos were randomly paired together to make it look as though they were having a real interaction. Observers perceived greater levels of synchrony in real interactions compared to the chance levels present in the pseudo interactions. This suggested that real interactions are characterised by synchronised behaviour and this can be perceived by untrained observers. However, this approach only taps into the general level of coordination present, and there is some evidence the level of synchrony perceived by the observers is actually correlated with other social cues such as smiling (Cappella, 1981) or physical similarity (Lumsden et al., 2012).

More recent studies have used automatic techniques to record high-resolution time series data about people's motion during naturalistic conversations (Paxton & Dale, 2013). One technique is to use a frame-differencing algorithm to automatically assess differences in each person's body movements from one video frame to the next in order to calculate their overall movement or 'motion energy' (Fujiwara & Daibo, 2016; Paxton & Dale, 2013; Fabian Ramseyer & Tschacher, 2010; Schmidt, Morr, Fitzpatrick, & Richardson, 2012). Although this has the advantage that specialist equipment is not needed, there are currently few established frame-differencing methods (Paxton & Dale, 2013) and this approach lacks resolution to capture specific head or limb movements. In order to achieve more specific data from individual body parts, another option is to use motion capture technologies. These include 3D depth-sensing cameras such as Kinect (e.g. Won, Bailenson, Stathatos, & Dai, 2014), as well as wearable motion sensors that can be directly attached to the people interacting (Feese et al., 2011; Poppe et al., 2013). Overall, automatic behaviour recording techniques can provide more objective and detailed data about an interaction compared to video coding by trained researchers or untrained participants. However, a major challenge is then to analyse the level of interpersonal coordination in the data.

5.2.2 Using Wavelet Analysis to Measure Interpersonal Coordination

Wavelet analysis is a recent approach to measuring interpersonal coordination in motion data that has been used to evaluate free conversations (Fujiwara & Daibo, 2016), musical improvisation (Walton et al., 2015), and telling knock-knock jokes (Schmidt et al., 2014). Wavelet analysis is a type of spectrum analysis that transforms a time series into time-frequency space. When applied to one person's motion trace, this means that different frequencies or rhythms within their motion can be examined across the time course. *Cross-wavelet analysis* gives the degree of similarity between two time-frequency transforms describing movements from each person in a dyad. The similarity is calculated as a correlation and is often termed ‘cross-wavelet coherence’. In general, high coherence is interpreted as a high degree of synchronisation (Fujiwara & Daibo, 2016), because it indicates that two people are moving with the same frequency but not necessarily in matching forms. However, when high coherence between similar behaviours is observed (e.g. both people waving their arms or both nodding), this could also be interpreted as mimicry because the movements match in form as well as frequency. Perfect mimicry, such as being mirrored exactly by a virtual avatar, would lead to maximum coherence at the frequency corresponding to the mimicry time delay (e.g. at 0.5 Hz for a mimicry delay of 2s). Note that the cross-wavelet approach gives the amount of coherence for the dyad, rather than indicating how much one person mimicked or synchronised with the other. As well as the amount of coherence, cross-wavelet analysis also calculates the phase relationship between two people, which can indicate whether one person led or followed the other at a particular rhythmic frequency, and how long this leader-follower pattern was sustained over time. For primers and tutorials on using wavelet analysis in the context of social research, see Issartel et al. (2006) and Issartel, Bardainne, Gaillot & Marin (2015).

Wavelet analysis is still a very new tool for studying naturalistic social interactions, and there have only been a handful of studies using this approach. These include initial proof-of-concept studies using small sample sizes to demonstrate how cross-wavelet analysis can be applied to time series data about the head or limb movements of two people interacting together (Issartel et al., 2006; Varlet, Marin, Lagarde, & Bardy, 2011; Walton et al., 2015). In one such study, Issartel et al (2006) demonstrated that 6 pairs of participants were unable to avoid unintentionally coordinating their arm movements, even when instructed to do so, during a free movement task where they swung their arms while seated. Subsequent studies using larger sample sizes have shown that people with dance training are better at synchronising with other people or external rhythms (Sofianidis, Elliott, Wing, & Hatzitaki, 2014; Washburn et al., 2014) and have demonstrated that bodily synchronisation occurs within nested timescales during a scripted task (Schmidt et al., 2014).

Recently, one study has used wavelet analysis to test whether the levels of cross-wavelet coherence between two people having a naturalistic conversation are greater than chance levels obtained from pseudo interactions. Fujiwara & Daibo (2016) carried out cross-wavelet analysis on 31 dyads engaged in 6-minute conversations and compared the amount of motion coordination (averaged over all component frequencies) to the amount in pseudo interactions where members of different dyads were randomly paired together. They found significantly greater coordination in the real interactions, replicating earlier findings (Bernieri, 1988; Bernieri et al., 1988; Schmidt et al., 2014) and adding further evidence that people spontaneously coordinate their body movements in natural conversations. In order to further examine the pattern of how people coordinate, Fujiwara & Daibo split the data from the genuine interactions into different frequency bands. They found that coherence levels gradually declined across the frequency spectrum, from lower (less than 0.025 Hz) to higher (4 Hz) frequencies. These initial results suggest that interpersonal coherence may vary at different motion frequencies, with less coherence at high frequencies close to 4 Hz. However, the authors did not test whether this pattern was also present in pseudo interactions, so it is unclear whether the same pattern would emerge in randomly paired data or whether true interactions involve more coherence at low frequencies than high frequencies.

Wavelet analysis versus cross-recurrence analysis. An alternative approach to wavelet analysis is to use cross-recurrence analysis to assess the degree of similarity or coordination between time series from two people interacting. Cross-recurrence analysis is based on recurrence quantification analysis (RQA). RQA was designed in the 1990s to identify recurrent patterns of overlap within time-lagged series (Shockley, Butwill, Zbilut, & Webber, 2002; Zbilut, Giuliani, & Webber, 1998a, 1998b) and originally had its main applications in the biological and physical sciences (D. C. Richardson, Dale, & Shockley, 2008). RQA involves two key steps: the first is to identify ‘recurrences’ in the time series, when the system being measured revists a similar state that is has been in before. The second step is to quantify the number and nature of the recurrences. In straightforward RQA, these steps are carried out for one time series representing one system. In cross-recurrence analysis, time series from two different systems are analysed and the recurrence between them represents the degree of overlap between their trajectories (Shockley, Baker, Richardson, & Fowler, 2007; Shockley, Santana, & Fowler, 2003).

The application of cross-recurrence for measuring interpersonal coordination has been developed by Kevin Shockley and colleagues over the past two decades. In response to traditional research that relied on subjective coding to assess interpersonal coordination, they have aimed to use cross-recurrence analysis as a more objective approach which allows for objective quantification of the behavioural coupling between two or more people engaged in a dynamic social interaction (Shockley et al., 2007, 2002; Shockley, Richardson, & Dale, 2009; Shockley et al., 2003). The key variables derived from cross-recurrence analysis are the percentage recurrence (%REC), or degree of shared movements, and the maximum line of recurrence points (MAXLINE), which indicates the length of the longest sustained overlap between interaction partners.

In a key study demonstrating this approach, Shockley et al. (2003) investigated the coordination of two people's postural sway during a social interaction task. When standing upright, the body naturally sways in continuous movement to maintain balance, especially to adjust for the movements made while breathing, speaking and gesturing. In order to track postural sway, Shockley et al. attached motion sensors to the heads and hips of two participants who engaged in a conversation task. They were each given a similar cartoon picture and were asked to discuss the pictures in order to find out subtle differences between them. Participants could move and gesture naturally, but would either face each other or away from each other while discussing their picture with either the other particiapant or a confederate (although the other participant was always present). Shockley et al. found that when participants completed the task with each other, their postural sway showed significantly greater coordination in terms of both %REC and MAXLINE, compared to when they each interacted with a confederate. This effect was independent of whether the participants were looking at each other or not. However, coordination of head movements did not differ across any conditions. The results from this key study suggested that cross-recurrence analysis can measure the dynamics of coordination during true social interactions, and this was supported by a follow-up study (Shockley et al., 2007).

Richardson and Dale (2005) have also used cross-recurrence analysis to measure the timing of joint gaze between speakers and listeners. The ‘speaker’ participants’ voices and eye movements were recorded while they looked at a set of TV sitcom characters and spoke spontaneously about the TV show. Segments of their speech were played to the ‘listener’ participants, whose eye movements were also tracked while they looked at the same display of sitcom characters. The timing of joint gaze between speakers and listeners was assessed by comparing the amount of cross-recurrence between eye positions at successive time lags. They found that a listener was most likely to be looking at the same sitcom character as the speaker after a lag of approximately two seconds. In subsequent studies, Richardson, Dale and Kirkham (Daniel C. Richardson, Dale, & Kirkham, 2007) used cross-recurrence analysis to investigate gaze coordination during live interactions.

Wavelet analsysis and cross-recurrence analysis may offer different advantages and disadvantages for investigating patterns of spontaneous coordination between two people’s movements, although the two techniques have not been directly tested against each other. When it comes to *quantifying* coordinated movements, cross-recurrence analysis and wavelet analysis are both suitable for studying natural social interactions, because neither of them make assumptions about the stationarity of data (Issartel et al., 2015; Shockley et al., 2007). Cross-recurrence analysis may also have the additional strength that it does not make assumptions about the distribution of data (Shockley et al., 2007), whereas wavelet analysis fits the data against a particular wavelet function and assumes that it has wavelet-like properties. This could mean that cross-recurrence analysis is a more versatile analysis tool than wavelet analysis. When it comes to investigating the *timing parameters* of interpersonal coordination, cross-recurrence analysis and wavelet analysis offer different kinds of insight: Cross-recurrence analysis has been used to assess coordination at successive time lags (Daniel C. Richardson & Dale, 2005), which could be seen as a more sophisticated alternative to the traditional cross-correlation of time series. In contrast, wavelet analysis allows the assessment of coordination at each component frequency across the whole spectrum that makes up a given motion trajectory. In other words, wavelet analysis could be used to reveal the rhythmic tempi at which motion is coordinated in a dyad, rather than the typical time delay between one person matching the other's movements. Neither of these analysis options is inherently better or worse than the other, but in this chapter we chose to apply wavelet analysis in order to investigate the rhythmic characteristics of mimicry in fine detail.

5.2.3 Interpersonal Coordination at Different Frequencies

A major advantage of wavelet analysis is that it can be used to measure interpersonal coordination at different frequencies, and could therefore provide insight into the natural rhythmic characteristics of mimicry. At present this is something we know very little about. It has been suggested that mimicry occurs with natural delays of 2-5 seconds (Leander et al., 2012; van Baaren, Decety, et al., 2009), based on anecdotal evidence (although estimates vary from 0 to 10 seconds; Chartrand & Bargh, 1999; Stel et al., 2009). This would correspond to movement coherence within a frequency band of roughly 0.5 – 0.2 Hz, which could be a possible range where we should expect mimicry to occur. However, this very rough estimate is based on our beliefs about easily observable mimicry behaviours such as body posture shifts and gestures. Coordination could also occur between many other subtle behaviours present in a conversation. For example, interpersonal coordination has previously been linked to the natural rhythm of speech (Condon & Ogston, 1966; Hadar, Steiner, Grant, & Rose, 1983a, 1983b; Kendon, 1970), which is around 5 Hz (Morrill, Paukner, Ferrari, & Ghazanfar, 2012; Ohala, 1975).

Further research suggests that facial expressions at frequencies of approximately 2 - 7 Hz are also important for the interpretation of speech (Chandrasekaran, Trubanova, Stillittano, Caplier, & Ghazanfar, 2009). Other studies have demonstrated entrainment in covert rhythms such as heart rate (1 – 1.7 Hz; Konvalinka et al., 2011) and breathing (0.2 - 0.3 Hz; Pellegrini & Ciceri, 2012; Warner, 1996). Human body movement is also thought to be somehow tuned to frequencies around 2 Hz based on musical preference and recall of these tempi (Noorden & Moelants, 1999). While it is not possible here to review all of the rhythmic behaviours that could be present in a conversation, Table 5-1 lists some common rhythmic behaviours in order to illustrate a range of motion frequencies from around 0 – 5 Hz which might be relevant for interpersonal coordination. Without any previous systematic tests of interpersonal coordination for different frequencies, it is hard to make firm predictions about the frequencies at which people naturally coordinate. However, the behaviours reviewed in Table 1 suggest natural head and body coordination is likely to occur at frequencies lower than speech production at 5 Hz (c.f. Fujiwara & Daibo, 2016). Therefore, in the present study we examine interpersonal coordination across a full frequency spectrum from 0 to 5 Hz.

Table 5-1. Motion frequencies of common rhythmic behaviours.

Frequency	Rhythmic behaviour	Reference
0.2 – 0.3	Breathing while speaking	(McFarland, 2001)
0.43	Blinking	(Bentivoglio et al., 1997)
1 – 1.7	Heart beat	(Jose & Collison, 1970)
1.8 – 3.7	Ordinary head motion in conversation	(Hadar et al., 1983b)
2	Average beat to music	(Noorden & Moelants, 1999)
2	Walking	(MacDougall & Moore, 2005)
2 – 7	Facial expression accompanying speech	(Chandrasekaran et al., 2009)
3.7 +	Fast head motion in conversation	(Hadar et al., 1983b)
5	Speech	(Morrill et al., 2012; Ohala, 1975)
5	Laughter	(Luschei, Ramig, Finnegan, Baker, & Smith, 2006)

5.2.4 The Present Study

In the present study we recorded movements from participant dyads engaged in a photo description task. This task was based on previous studies (Chapters 3 and 4), although other research has also used a very similar task to study body coordination with cross-recurrence analysis (Shockley, Santana, & Fowler, 2003). In our task, participants took turns to describe photos to each other for 90 seconds. So that we could explore whether there were any differences in interpersonal coordination between monologue and dialogue interactions, each turn was split into two parts. For the first 30 seconds the speaker described the photo in monologue and for the remaining 60 seconds the listener could ask questions and both participants could converse together. We also included a scrambled sentences priming task which has been shown to subliminally increase or decrease levels of automatic imitation in finger-tapping paradigms. For half of the trials, participants were primed prosocially and for the other half they were primed antisocially. We recorded the position and rotation of motion sensors on each participant's head and torso, at a rate of 60 Hz. Their conversation was also video and audio recorded, although we do not analyse the video or audio data in this chapter. From the motion data we primarily examined four head signals: head motion energy, head yaw (turning) head pitch (nodding), and head roll (tilting). We used wavelet analysis to calculate levels of interpersonal coordination in real trials and in pseudo trials, which were created by matching data from different trials within the same pair.

As wavelet analysis is a relatively new technique in social cognitive research, the present study consisted of a pilot phase and a final phase. The pilot phase was highly exploratory, whereas the final phase reported here aimed to rigorously replicate our pilot results following a strict pre-registered analysis pathway (Hale & Hamilton, 2016a). In both phases we used the same method of data collection. In the pilot phase, a student supervised by Joanna Hale (Francesco Buccheri) collected a dataset from 20 dyads. We then carried out exploratory wavelet analyses on multiple head and torso signals from the raw pilot data (Hale & Hamilton, 2016a). Surprisingly, the pilot results suggested that real interactions involved *less* cross-wavelet coherence in head movements than pseudo interactions, at movement frequencies of approximately 1-7 Hz. The results also suggested that prosocial priming might increase head coordination in monologue interactions. However, we carried out many exploratory tests and it is possible these results could reflect false positives or spurious trends. Therefore, we aimed to carry out more rigorous tests on a new dataset. A different research assistant from our lab collected the final dataset, following the same data collection procedure as the pilot phase. Based on our pilot analyses, we finalised the analysis pathway for the final dataset and pre-registered this pathway along with Matlab scripts for performing the analyses. The pre-registered analysis pathway focused only on head movements, where we saw the strongest results in the pilot data. In this chapter, we also report additional analyses examining coherence in torso movements.

5.3 Method

5.3.1 Participants

All participants were recruited through email advertisements to a departmental database of research volunteers. We recruited 31 dyads in total (M_{age} = 22.3 years, SD_{age} = 2.9 years). Participants were paired up as dyads on the basis of their availability and preference for the same time slot. All participants gave written consent and were paid £7.50 for 1 hour. The study received ethical approval from the UCL Research Ethics Committee (Application 5713/001) and was performed in accordance with the 1964 Declaration of Helsinki.

5.3.2 Equipment

Motion tracking system. We used a Polhemus magnetic motion tracking device (Polhemus Inc., Vermont) to detect participants' movements (Figure 5-1A). This device detects the position and rotation of sensors within a low-intensity electromagnetic field, with a specified frequency (we used 60 Hz). Participants were fitted with two sensors.

One sensor was fixed on the participant’s forehead using a Velcro cloth band. Another sensor was attached to the participant’s upper back using microporous tape.

Audio and video recording. Both participants wore a lapel microphone and their voices were recorded. A video camera mounted on a tripod recorded the session, offering a clear view of the participants’ seated bodies. We do not report any audio or video data in this chapter.

Laboratory set-up. The laboratory set-up is illustrated in Figure 5-1A. In the laboratory there were two wooden stools for the participants, facing each other at a distance of approximately 1.5m. Between the stools there was a wooden structure to hold the Polhemus transmitter device which generated an electromagnetic field of approximately 2m diameter around the participants. A projector screen to the side of the participants showed instructions throughout the session, and audio speakers above the projector provided audio cues. A curtain separated participants from the experimenter, who remained in the room but did not interact with participants during the experiment and could not be seen.

Data recording software. We used Vizard software to display instructions on the projector screen and trigger audio cues at the correct times. Data from the motion tracker was read into Vizard at a rate of 60 Hz. We used Vizard to save time-stamped motion data, labelled with the current experimental conditions and block and trial numbers.

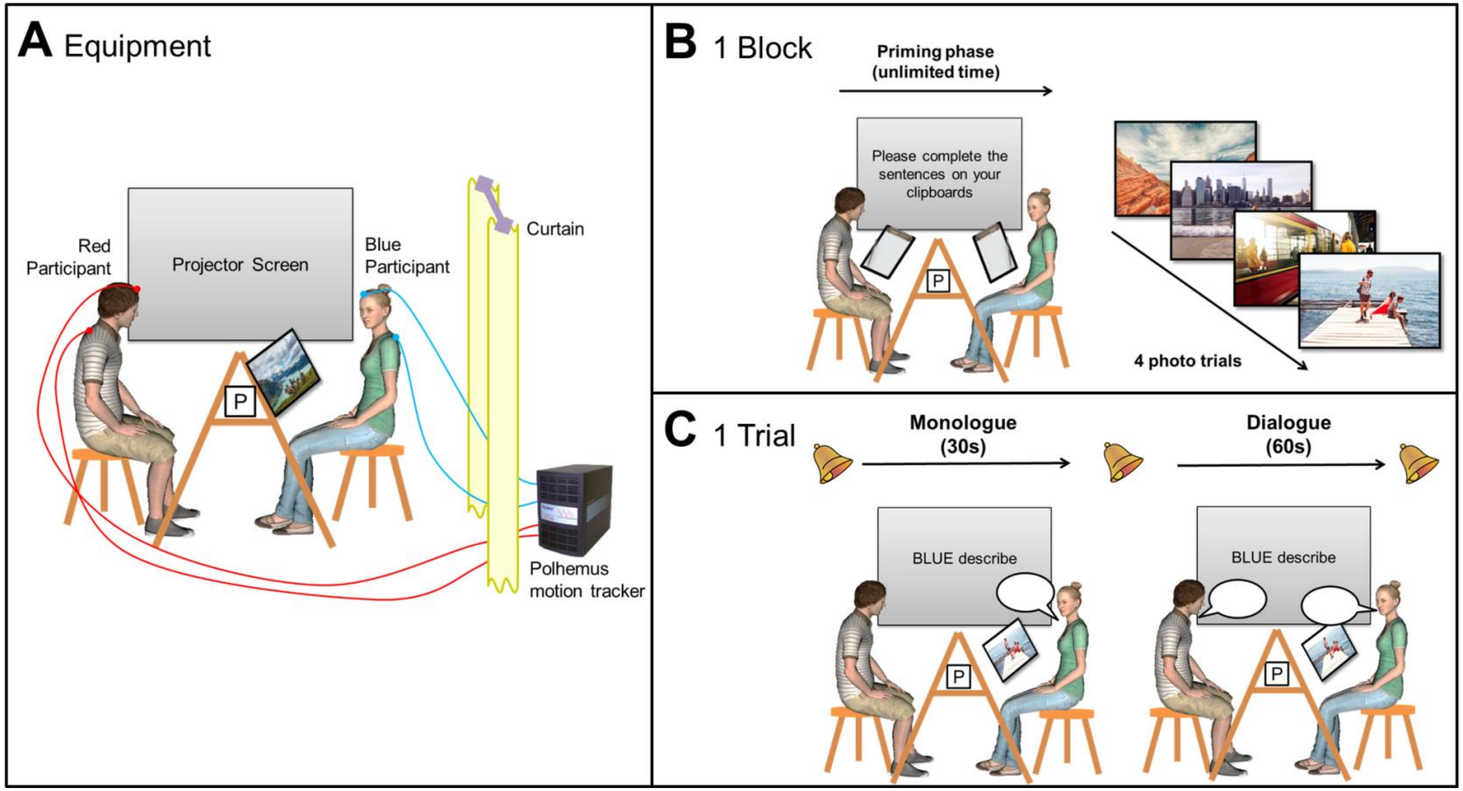

Figure 5-1*. Equipment and procedure.* Panel A shows the equipment and lab arrangement. Red and Blue participants were seated facing each other. Each participant wore a motion tracking sensor on their forehead and upper back. Between them a wooden frame held the transmitter device for the Polhemus motion tracker. Instructions were displayed on a projector screen and participants were separated from the experimenter by a curtain. Panel B shows one block of the experiment. At the start of the block, participants completed a scrambled sentence priming task on clipboards. The priming task was followed by four photo description trials. Panel C shows one trial. One participant (e.g. Blue) was the speaker for the trial. The speaker described a photo in monologue for 30 seconds, then both participants could speak in dialogue for 90 seconds. The start of the trial, the end of the monologue section and the end of the trial were signalled with audible beeps.

5.3.3 Procedure

Dominic Oliver carried out the data collection. Two participants were scheduled to arrive at the laboratory at the same time, and were met by the experimenter. After giving informed consent to take part, the participants were fitted with motion sensors and microphones. They were randomly assigned to be the 'Red' participant or 'Blue' participant. These labels were used to distinguish dyad members during the experiment and in the recorded data.

The experiment was split into one practice block and four experimental blocks. During the practice block, the experimenter gave the participants verbal instructions about each task, in addition to instructions which appeared on the projector screen. During the experimental blocks the participants only received instructions on the projector screen. Each block began with a scrambled sentences priming task, followed by a photo conversation task. The experimental blocks alternated between prosocial and antisocial priming conditions. The order was counterbalanced across participants in a randomised fashion.

Scrambled sentences priming task. Although it is not the main focus of the present chapter, we included a task to investigate whether social priming would affect levels of coordination between dyads. At the start of each experimental block, participants individually completed a priming task taken from Wang & Hamilton (2013) using paper and pen (Figure 5-1B). The task was to make a correct sentence using six words out of a jumbled list of seven words, where two words were already in the correct order. When completed, each sentence described an interaction that was either prosocial (e.g., "*Alex and Zoe enjoy their holiday in Hawaii*") or antisocial (e.g., "*Stuart and Eva fight over the last biscuit*"), and all primes had a third person perspective. In the practice block, the sentences presented neutral factual information instead (e.g., "*The Nobel prize ceremony is held in Sweden*"). Individual participants completed different sentences to each other but were both primed with the same condition (e.g. both prosocial).

Picture description task. Following the priming task, participants completed a picture description task (Figure 5-1C). Stimuli for the task were colour photos showing human figures, taken from the National Geographic website and printed on heavy card. In order to generate a conversation between participants, they were asked to take turns at describing a photo to each other (Chartrand & Bargh, 1999). Each turn describing a photo corresponded to one trial. There were two trials in the practice block and four trials per experimental block (see Table 5-2 for an example trial sequence). Before the start of each trial, participants were instructed who should be the speaker (e.g. Red participant) and listener (e.g. Blue participant) for that trial. The trial was split into two parts. In the first part (monologue), the speaker was instructed to describe the photo for 30 seconds and the listener was instructed not to speak. In the second part (dialogue), the speaker and listener were instructed to converse freely for 60 seconds, and the listener was instructed they could ask questions about the photo. Audible beeps indicated the start of the trial, the start of the dialogue section and the end of the trial. A timer on the projector screen also counted down the time left in the monologue and dialogue sections.

Interaction quality questionnaire. At the end of the study, participants individually completed a questionnaire about the quality of their interaction. The questionnaire data is not analysed or reported here. Our virtual mimicry studies (Chapters 3 and 4) and piloting of the present study showed weak or null relationships between mimicry and self-reported rapport, trust and other aspects of interaction quality. Therefore the present pre-registered study did not aim to investigate any social consequences of mimicry.

Finally, participants were asked to write down what they thought was the purpose of the study, and were debriefed and paid by the experimenter.

5.3.4 Contributions

Francesco Buccheri carried out pilot data collection, supervised by Joanna Hale. Dominic Oliver carried out participant recruitment and data collection for the present study, following training by Joanna Hale. Joanna Hale carried out all data pre-processing and analyses in MATLAB, which are described below.

Table 5-2. Example trial sequence.

Block	Priming condition	Trial	Speaker
1	Prosocial	1	Red
		2	Blue
		3	Red
		4	Blue
2	Antisocial	5	Red
		6	Blue
		7	Red
		8	Blue
3	Prosocial	9	Red
		10	Blue
		11	Red
		12	Blue
4	Antisocial	13	Red
		14	Blue
		15	Red
		16	Blue

Note. The order of prosocial and antisocial priming conditions was counterbalanced across dyads.

5.4 Analyses

5.4.1 Data Exclusion Criteria

We excluded data from dyads who met any of the following criteria:

1. Participants knew each other before the study
2. Motion data was not recorded due to technical failure of the equipment or task software
3. Motion sensor(s) moved or fell off during the study
4. More than 50% of their data was missing or not suitable for wavelet analysis

Before carrying out any analyses, data were excluded from 5 dyads who met one of our exclusion criteria. For 1 dyad, participants knew each other before the study. For 2 dyads, motion sensors moved or fell off during the study. For 2 dyads, more than 50% of their data was missing. The final sample consisted of 26 dyads (M_{age} = 22.3 years, SD_{age} = 2.9 years). There were 16 same-gender dyads and 10 mixed-gender dyads (34 female and 18 male participants). When carrying out the analyses, we also excluded individual trials if wavelet analysis could not be performed (e.g. this could happen if there is a very large jerk or jump in the motion data).

5.4.2 Data Format

For each participant, motion was recorded from two sensors. One sensor was on the head and the other was on the torso (Figure 5-1A). Each sensor gives three data channels specifying its position (x, y and z) and three channels specifying its rotation (yaw, pitch and roll). The yaw, pitch and roll signals roughly correspond to head turning, nodding and tilting. However, the rotations are calculated relative to a fixed frame of reference and the calculations are performed in a fixed order (Figure 5-2). This means that the yaw reading affects the pitch reading (Figure 5-2C), and both the yaw and pitch readings affect the roll reading (Figure 5-2D). Thus, pitch and roll do not perfectly correspond to the head nodding and tilting angles we would calculate if we were working with the head as our local frame of reference. However, because people do not usually make very large rotations of the head, we interpret the pitch and roll signals as close approximations of head nodding and head tilting, respectively.

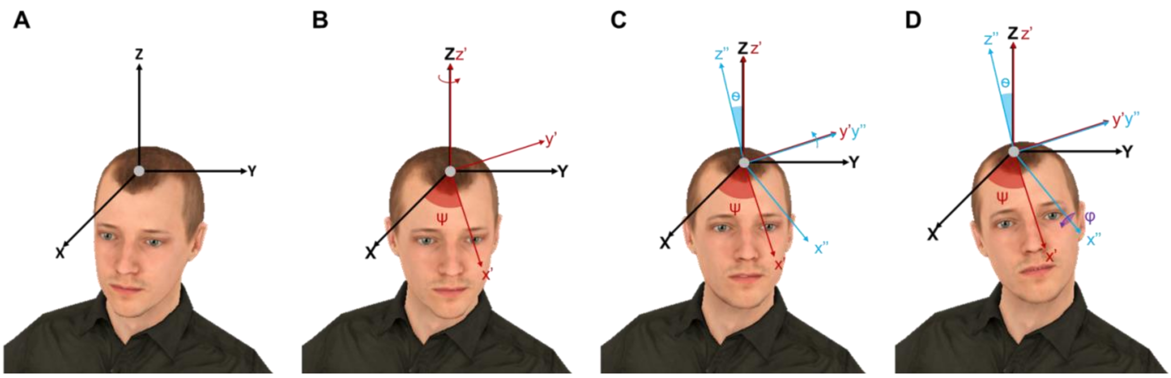

Figure 5-2. Rotation transforms for calculating yaw, pitch and roll signals. The rotations are calculated relative to a fixed reference frame, axes XYZ (Panel A). Three transforms give the final yaw, pitch and roll rotations. The first transform is a rotation about the Z axis, which gives the yaw angle, ψ (Panel B). The second transform is a rotation about the y’ axis, which gives the pitch angle, ɵ (Panel C). The third transform is a rotation about the x’’ axis, which gives the roll angle, φ (Panel D).

We recorded data from the channels above at a rate of 60 data points per second. One trial lasted 90 seconds, split into 30 seconds of monologue and 60 seconds of dialogue (Figure 5-1C). Thus, the complete dataset for a single trial has 24 channels: 3 head position, 3 head rotation, 3 torso position and 3 torso rotation for each participant, over 90 seconds (approximately 5400 data points per channel). There were 16 trials (separate time series) per dyad. Table 5-2 gives the experimental conditions of each trial. We read the raw data for each dyad into Matlab as 16 x 5400 matrices.

5.4.3 Pre-Processing

We carried out some pre-processing on the raw data. First, we trimmed the data by discarding the first 100 time points (1.7s of the trial) and all time points after the 5,250th point (87.5s into the trial; note that we originally specified 5300 in our pre-registration but some trials were shorter than 5300 data points). This ensured that all time series are of exactly equal length, which was necessary in order to average across trials and participants. It also avoided the inclusion of any unusual or jerky movements made at the start or end of the trial (signalled by audible beeps), such as turning to look at the projector screen. Then we corrected for circularity in the rotation data channels, to deal with cases where a change in orientation from -355° to 5° appears to Matlab like a large change rather than only a 10° movement past zero. We used a Matlab script which detects changes in angle of more than 270° and then adds / subtracts 360° to the data after the jump, in order to end up with continuous data. Next, we de-trended each data channel by subtracting the mean value, and applied a 7th order Butterworth low pass filter with cut-off frequency of 0.9 to reduce noise in each data channel. Finally, we applied a timing correction to deal with any tiny inaccuracies in the data time stamps which are incompatible with wavelet analysis. We created a timeline with the same number of data points linearly spaced from the first to last data point, and replaced the recorded timeline with this precisely equidistant one. Note we did not exclude any data for being noisy or poor quality, but rather kept all raw data from all channels at this stage.

5.4.4 Main Signals

We analysed 4 main signals. First, we calculated head motion energy as the sum of the motion velocities in x, y and z directions. Second, we kept head yaw, head pitch and head roll data as separate signals. This allowed us to preserve information about how our participants moved in different ways, such as nodding the head versus tilting the head from side to side. These different movements might provide different social cues in an interaction and show different patterns of coordination.

5.4.5 Analysis of Signal Power

As an initial analysis, we examined the power spectral density (PSD) estimate for all signals using Welch's method. First, for each participant and each trial we calculated the PSD estimate. Then, separately for each experimental condition, we averaged together the PSD estimates from participants in the same dyad across all trials. Finally, we carried out t-tests to determine whether there were significant differences in power between prosocial and antisocial conditions, and between speaker and listener conditions. This allowed us to determine if there were overall differences in participants' movement behaviour between the different priming conditions and speaker/listener roles.

5.4.6 Wavelet Analysis

We carried out wavelet analysis on each of our 4 signals, using the following pipeline (summarised in Figure 5-3). All wavelet analyses were carried out using the Matlab toolbox from Grinsted et al. (2004) with default parameters. First, we took two time series, representing 86s of social interaction from one trial between two participants. We calculated the wavelet transform of each time series, the cross-wavelet transform of the two time series together, and the cross-wavelet coherence in the interaction (Figure 5-3). We did this for all 16 trials, giving us 16 cross-wavelet coherence plots per dyad. Pilot analyses suggested that in a small minority of trials (mainly those with a single very jerky movement), the wavelet toolbox is unable to calculate the wavelet transform. Such trials were excluded from all analyses, and reported as missing data.

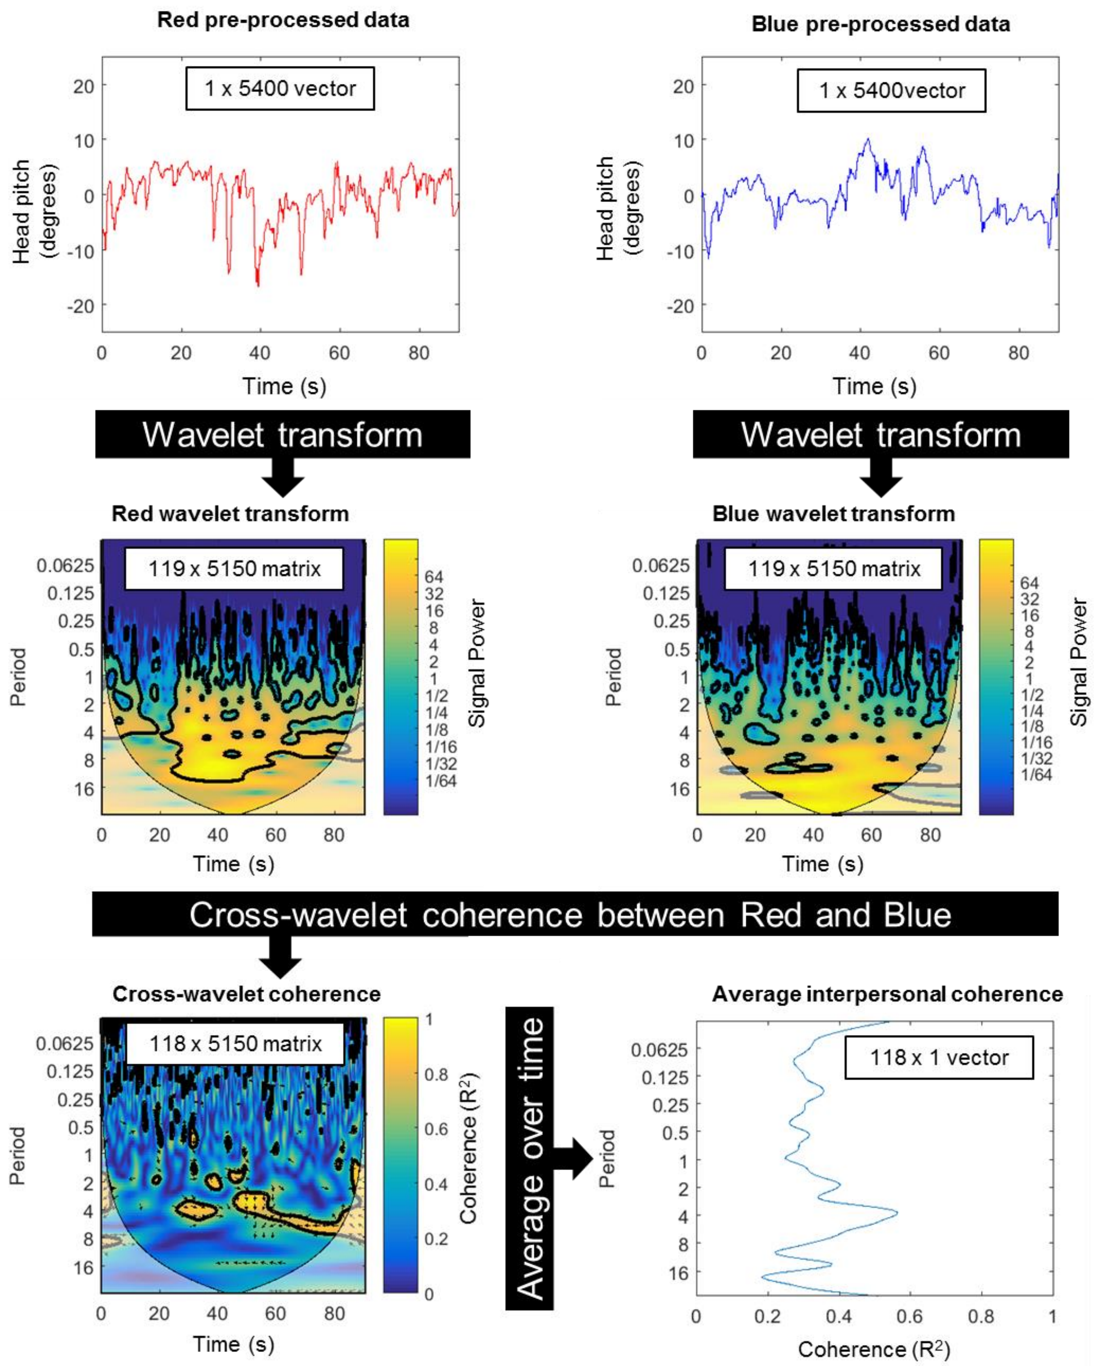

Figure 5-3. Wavelet analysis pipeline for one trial.

Next, we collapsed the cross-wavelet coherence over the time course of each trial. This is because we are interested in the periodicity of any interpersonal coherence, but not the specific time at which it occurred. We averaged the cross-wavelet coherence (R^2) at each periodicity over the whole trial (86s). We also separately average over monologue and dialogue sections of the trial, excluding 100 time points (1.7s) either side of the transition from monologue to dialogue, when there was an audible beep that might have triggered unusual head movements. Thus, 26.6s of monologue and 55.8s of dialogue were analysed from each trial.

The final output of the wavelet coherence analysis was a 1 x 118 vector of data points (one point for each wavelet periodicity from 0.03s to 31.4s) for each trial, representing the average wavelet coherence at each periodicity over that trial. We calculated these vectors for all 4 signals (head motion energy; head yaw; head pitch; head roll) and refer to these as the *coherence values* for each signal. For the following analyses, we converted periods to frequencies (period = 1/frequency). We also truncated the wavelet output to frequencies from 0 – 5 Hz, resulting in a 1 x 89 vector of coherence values. While our original pre-registration specified we would test all possible frequencies in our dataset (0-30 Hz), very high frequencies above 20 Hz could have been contaminated by dropped frames (affecting 2.7% of data points). In addition, frequencies above 5 Hz are unlikely to show meaningful interpersonal coordination (Fujiwara & Daibo, 2016).

5.4.7 Interpersonal Coherence in Real vs. Pseudo Interactions

A key test of interpersonal coordination is to compare coherence in real trials, where the two datasets entered come from the same interaction, with coherence in pseudo trials where two datasets from different interactions are entered into the algorithm (Bernieri & Rosenthal, 1991; Fujiwara & Daibo, 2016). Previous studies using this approach created pseudo trials by mixing datasets from different participants. We used a stricter approach where we mixed datasets from different trials within the same dyad and the same priming condition. Thus, our pseudo trials had the same general movement characteristics (e.g. overall signal power) as our real trials, and differed only in that the real trials represent a genuine live social interaction. To create pseudo trials, we matched up the Red participant's data from one trial with their Blue partner's data from a different trial. We only matched trials that involved the same priming condition and the same person speaking (Table 5-3). Note that for each real trial, there are 3 pseudo trial combinations, giving 48 pseudo trials per dyad. We carried out wavelet analysis on the pseudo dataset using the same pipeline as above. This gave a set of coherence values for each real trial of each dyad and each pseudo trial of each dyad.

Table 5-3. Example of generating pseudotrials for a particular dyad and trial order.

Block	Priming condition	Trial	Speaker	Trials with same conditions	True match		Pseudo match 1		Pseudo match 2		Pseudo match 3	
					R	B	R	B	R	B	R	B
1	Prosocial	1	Red	3, 9, 11	1	1	1	3	1	9	1	11
		2	Blue	4, 10, 12	2	2	2	4	2	10	2	12
		3	Red	1, 9, 11	3	3	3	1	3	9	3	11
		4	Blue	2, 10, 12	4	4	4	2	4	10	4	12
2	Antisocial	5	Red	7, 13, 15	5	5	5	7	5	13	5	15
		6	Blue	8, 14, 16	6	6	6	8	6	14	6	16
		7	Red	5, 13, 15	7	7	7	5	7	13	7	15
		8	Blue	6, 14, 16	8	8	8	6	8	14	8	16
3	Prosocial	9	Red	1, 3, 11	9	9	9	1	9	3	9	11
		10	Blue	2, 4, 12	10	10	10	2	10	4	10	12
		11	Red	1, 3, 9	11	11	11	1	11	3	11	9
		12	Blue	2, 4, 10	12	12	12	2	12	4	12	10
4	Antisocial	13	Red	5, 7, 15	13	13	13	5	13	7	13	15
		14	Blue	6, 8, 16	14	14	14	6	14	8	14	16
		15	Red	5, 7, 13	15	15	15	5	15	7	15	13
		16	Blue	6, 8, 14	16	16	16	6	16	8	16	14

Note. The order of prosocial and antisocial priming conditions was counterbalanced across dyads.

We carried out the following analysis three times, looking separately at monologue data, dialogue data and full trial data (monologue plus dialogue). Separately for the real dataset and the pseudo dataset, we averaged the coherence values across all trials for all dyads. This gave us group levels of coherence in real interactions and pseudo interactions. We then calculated a coherence difference for each dyad, representing the average coherence in real interactions minus the average coherence in pseudo interactions. Then we carried out t-tests on the coherence differences at each frequency

(89 tests). To correct for multiple comparisons, we used a false detection rate (Benjamini & Hochberg, 1995) in Matlab.

5.4.8 Effect of Prosocial and Antisocial Primes

We also carried out the following analysis separately for monologue data, dialogue data and full trial data: To test whether prosocial and antisocial primes lead to different levels of coordination, we calculated a coherence difference for each dyad, representing the average coherence in prosocial trials minus the average coherence in antisocial trials. We carried out t-tests on the coherence differences at each frequency, with p values corrected for multiple comparisons.

5.5 Results

5.5.1 Signal Power

We carried out t-tests to determine whether there were significant differences in the power spectral density (PSD) estimate of each signal between prosocial and antisocial conditions, and between speaker and listener conditions. Speakers showed significantly greater power than listeners for head yaw, head roll signals and head motion energy signals (Figure 5-4). They showed fewer differences in the head pitch signal, although speakers had significantly greater power at some frequencies. There were no significant differences in signal power between the prosocial and antisocial conditions in any signal (Figure 5-5).

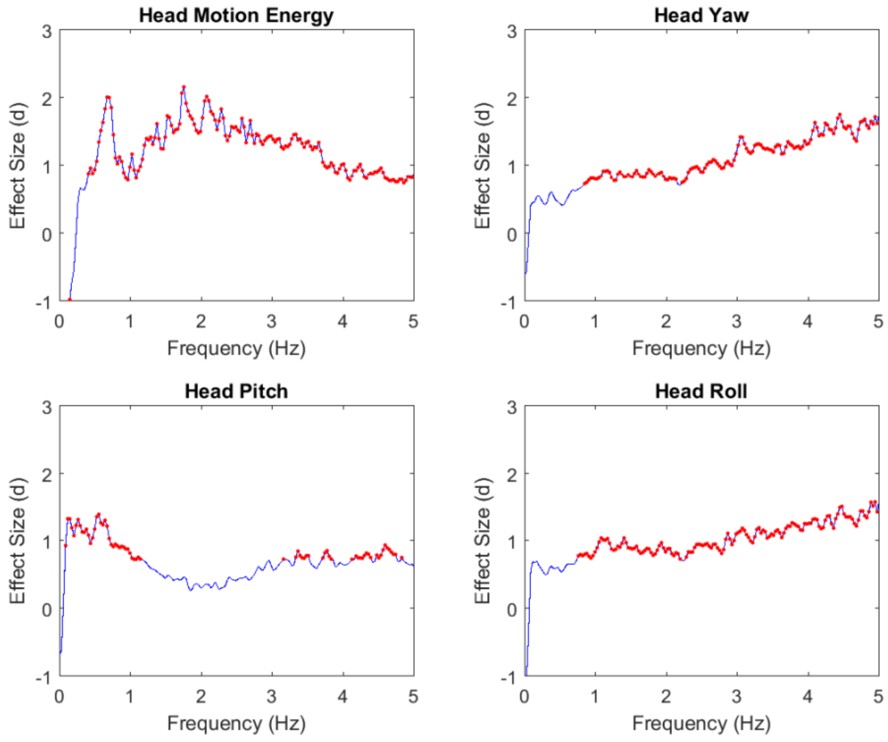

Figure 5-4. Difference in PSD estimates for speaker and listener head motion. Red dots indicate significant differences.

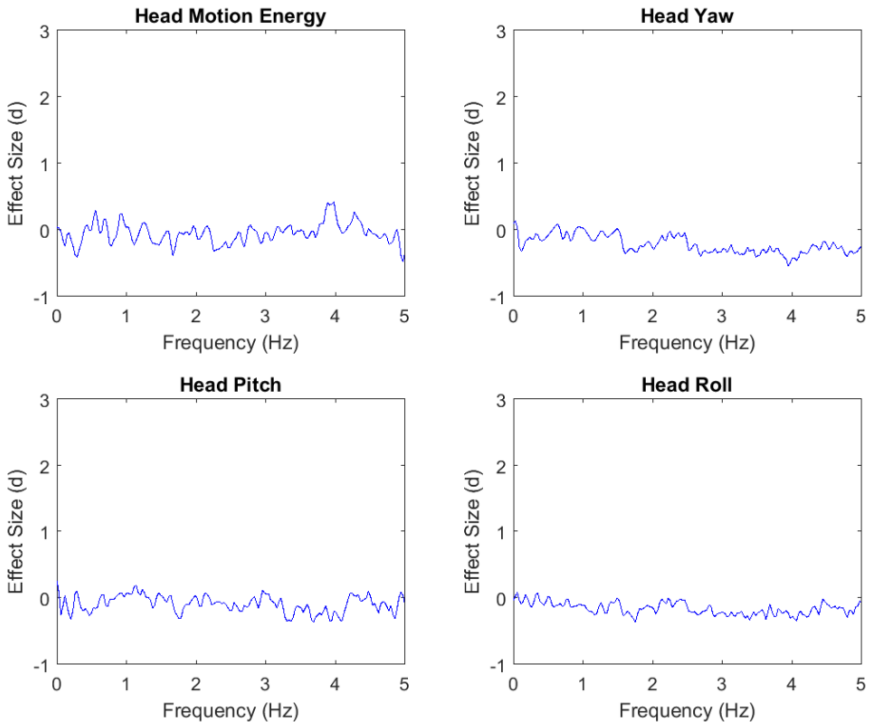

Figure 5-5. Difference in PSD of dyadic head motion in prosocial and antisocial conditions.

5.5.2 Cross-wavelet Coherence of Head Movements in Real vs. Pseudo interactions

Results are shown separately for each head signal in Figure 5-6 to 5-9. In each figure, panel A shows the coherence values for real and pseudo interactions. To assess the difference in cross-wavelet coherence between real and pseudo interactions, we performed t-tests at each frequency (89 tests). Panel B shows the results for full trials (including both monologue and dialogue sections). There were virtually no differences in results between the whole trial, the monologue section and the dialogue section, and therefore we only report the results for full trials. The plot in panel B shows the effect size (Cohen's *d*) for the difference between real and pseudo interactions. Red dots indicate frequencies where there was a significant difference between real and pseudo levels of coherence, with correction for multiple comparisons. Blue dots indicate uncorrected significant differences. Panel C shows the average signal power for the speaker and listener. Panel D shows the average phase relationship between speaker and listener as a rose plot from -180 to 180 degrees offset, where 0 degrees offset is perfect in-phase coordination.

For all four signals, there was a similar pattern of coherence across frequencies. Visually, this pattern could be split into two frequency ranges, above and below 1.5 Hz. Above 1.5 Hz, all rotation signals surprisingly showed significantly *less* coherence in real interactions compared to pseudo interactions (Figures 5-7 to 5-9). The motion energy signal also showed less coherence, but this was not significant after FDR correction (Figure 5-6). However, below 1.5 Hz all signals showed coherence close to or greater than chance levels during real interactions. Below 1.5 Hz, there were also some differences between the levels of coherence in each signal. Head pitch was significantly more coherent in real interactions than pseudo interactions at all frequencies from 0.07-1.1Hz (Figure 5-8), showing the strongest coordination above chance level of all the signals. Head yaw showed less coherence and was only significantly more coherent than chance at 0.09-0.11 Hz (Figure 5-7). Head roll (Figure 5-9) and motion energy (Figure

5-6) did not show any significant differences that met FDR correction at frequencies below 1.5 Hz. A visual inspection of the average phase relationship between speaker and listener suggests that coordination in the range below 1.5 Hz did not show any strong patterns of phase locking at a specific phase offset.

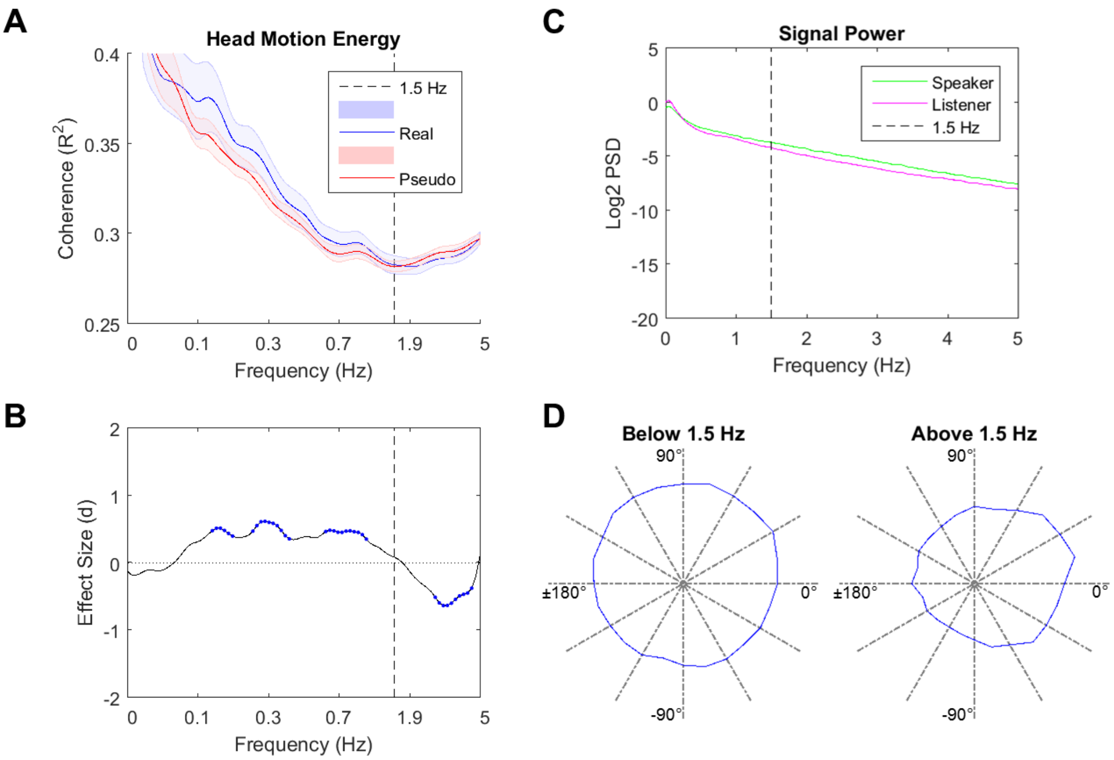

Figure 5-6. Head motion energy in real vs. pseudo interactions. Panel A shows the cross-wavelet coherence values for real and pseudo interactions. Panel B shows the effect size of the difference between real and pseudo interactions. Red dots indicate significant t-test results with FDR correction. Blue dots indicate uncorrected significant t-test results. Panel C shows the average signal power for the speaker and listener in real interactions. Panel D shows the average phase relationship between speaker and listener in real interactions.

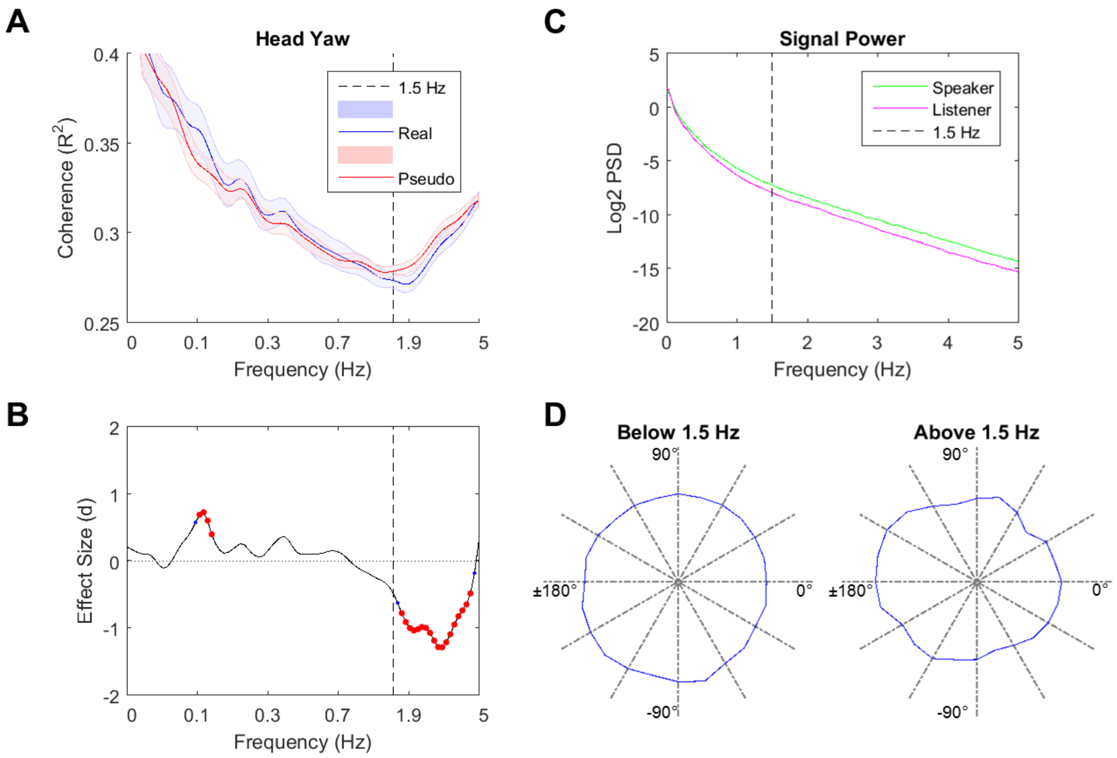

Figure 5-7. Head yaw in real vs. pseudo interactions.

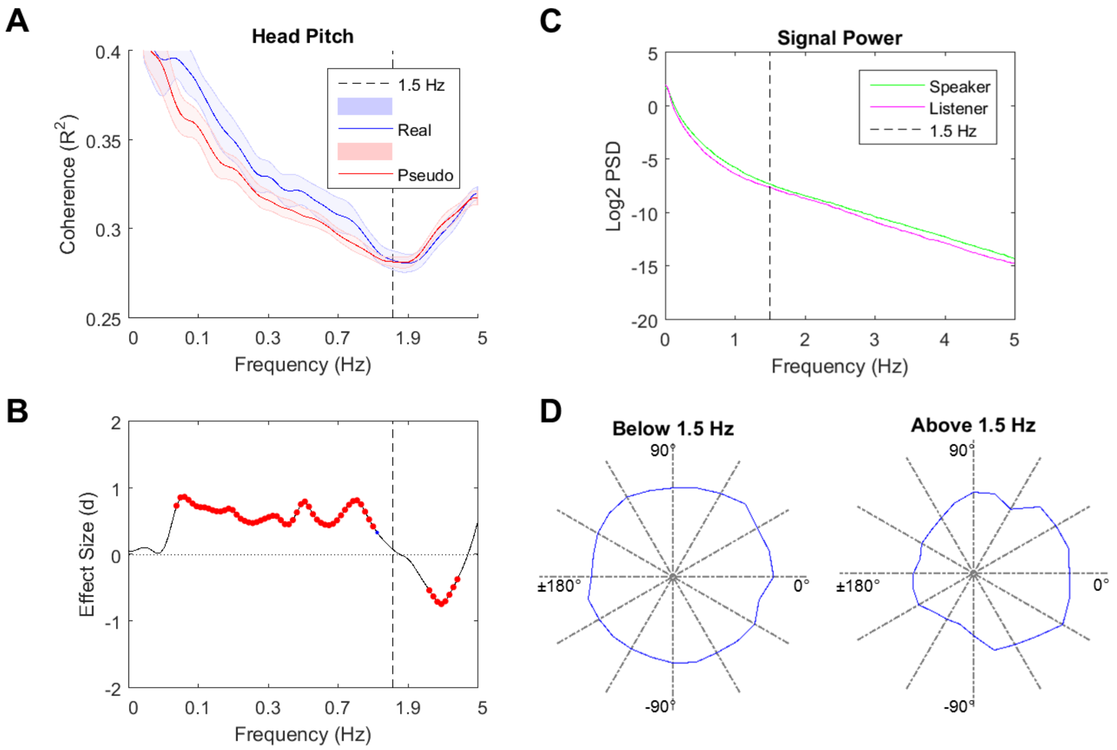

Figure 5-8. Head pitch in real vs. pseudo interactions.

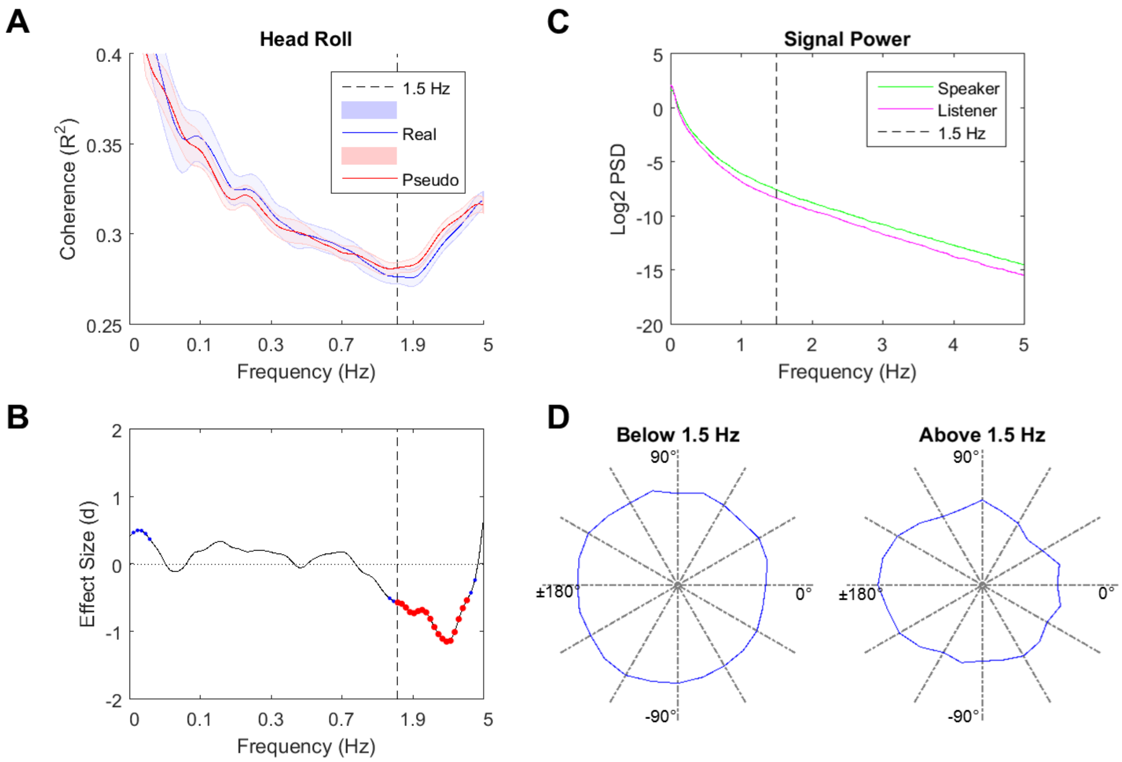

Figure 5-9. Head roll in real vs. pseudo interactions.

5.5.3 Effect of Prosocial and Antisocial Primes on Head Movements

The coherence values for prosocial and antisocial priming conditions are shown in Figure 5-9. To assess the coherence difference between trials following prosocial versus antisocial primes, we performed t-tests at each frequency. Figure 5-10 shows results for full trials (including both monologue and dialogue sections). There were no significant differences in coherence between prosocial and antisocial conditions that would meet a correction for multiple comparisons. The same results were found when we separately analysed the monologue and dialogue sections of the interaction.

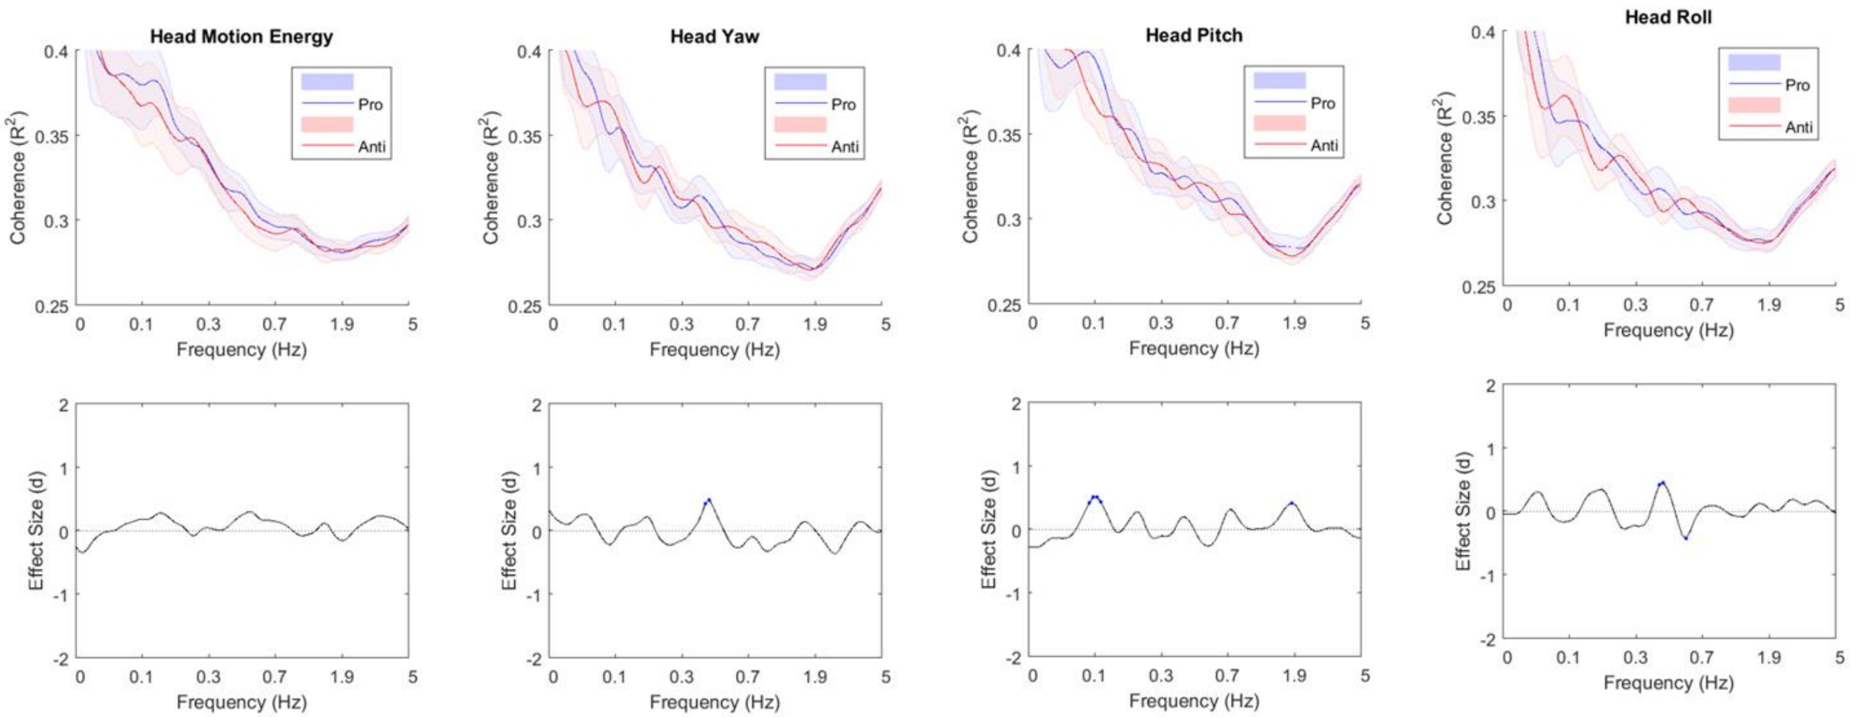

Figure 5-10. Coherence values in prosocial and antisocial conditions with effect sizes for the difference between conditions. Blue dots indicate uncorrected significant differences.

5.5.4 Additional Analysis of Torso Movements in Real vs. Pseudo interactions.

In the pre-registered analyses above we focused only on head movements. A natural extension was to apply the same analysis pathway to torso signals from the same dataset. In particular, we wanted to explore whether torso movements also show significantly less coordination in real interactions compared to pseudo interactions. Therefore, we calculated the cross-wavelet coherence for four torso signals using the same procedure as above. The torso coherence values for real and pseudo interactions are shown in Figure 5-11. To assess the difference in cross-wavelet coherence between real and pseudo interactions, we performed t-tests at each periodicity (as we did for head motion). Figure 5-11 shows the results for full trials. We did not separately examine monologue and dialogue sections of the trial, as these showed no differences in the head signals. As we found for head movements, at the frequency range below 1.5 Hz coherence in the torso signals was close to or greater than chance. There was significantly greater than chance coherence in torso motion energy at 0.13 - 0.22 Hz and 0.47 – 0.79 Hz, and in torso pitch at 0.13 – 0.16 Hz. However, at frequencies above 1.5 Hz no signals showed significantly less coherence in real interactions compared to pseudo interactions. In fact, at 3.6 - 5 Hz there was greater than chance coherence in motion energy of the torso during real interactions, and there was also significantly greater than chance coherence in torso pitch and roll at 4.8 – 5 Hz.

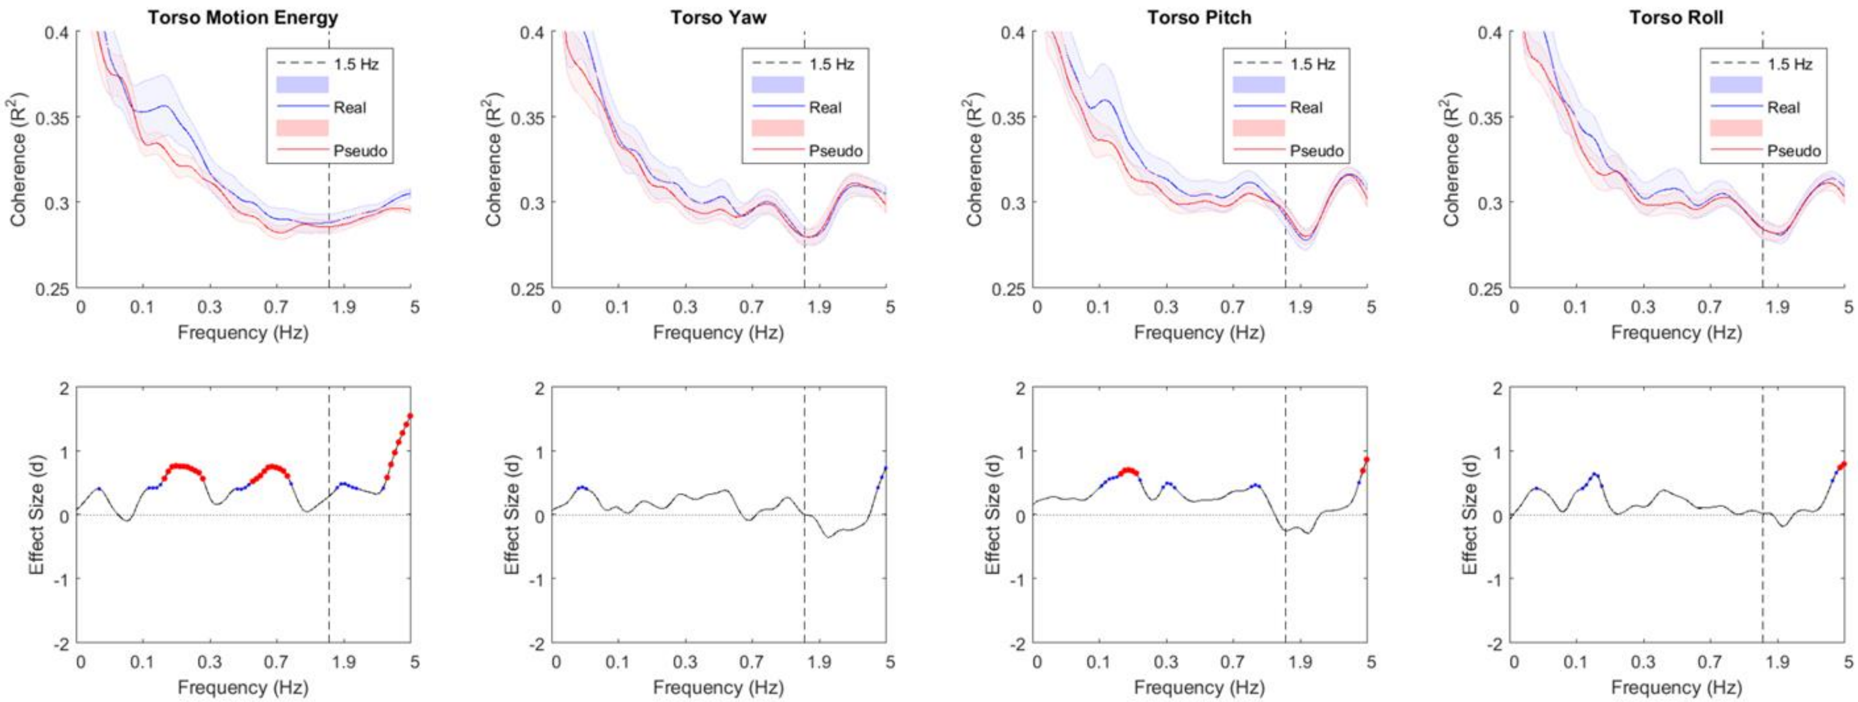

Figure 5-11. Coherence values for torso signals with effect sizes for the difference between real and pseudo interactions. Red dots indicate corrected significant differences. Blue dots indicate uncorrected significant differences.

5.6 Discussion

We used high-resolution motion capture to record the movements of participant dyads during a naturalistic conversation task with monologue and dialogue sections. In order to assess levels of interpersonal coordination we measured the cross-wavelet coherence of each dyad's head movements in terms of four signals: overall head motion energy, head yaw (turning the face), head pitch (nodding) and head roll (tilting side-to-side). We found that real interactions showed greater coherence than pseudo interactions, but this was only true for frequencies below 1.5 Hz and was strongest in the head pitch signal. Surprisingly, at higher frequencies between 1.5 and 5 Hz head movements showed significantly *less* coherence in real interactions compared to pseudo interactions (chance levels). No such pattern of divergence was seen for torso movements. In the coherence levels from the real interactions, we did not find any significant effects of priming using prosocial and antisocial scrambled sentences. In addition, all of our analyses showed practically identical patterns of coherence for both the monologue and dialogue parts of the conversation task. First we offer possible interpretations for the different patterns of coherence above and below 1.5 Hz, before discussing methodological implications of our data and future directions for research.

5.6.1 Decoupling of Head Movements at High Frequencies

It was a surprising result that all head rotation signals were significantly *less* coherent than chance levels during real social interactions. Dominant theories of interpersonal coordination would expect people to spontaneously coordinate (Bernieri & Rosenthal, 1991; Lakens et al., 2016; Marsh, Richardson, & Schmidt, 2009), but not to actively un-coordinate their movements. Our findings challenge these predictions by demonstrating that natural conversations are characterised by active decoupling of head movements at frequencies of around 1.5 to 5 Hz. These results do not simply suggest that dyads failed to coordinate their head movements in this range, but actually coordinated significantly less than we would expect from chance levels when data from within the same dyad was shuffled. Note that we observed the same pattern in our pilot dataset (Hale & Hamilton, 2016a). It is also worth noting that head movements in the frequency range of 1.5 – 5 Hz are fairly typical in conversations. Hadar et al. (1983b) recorded head rotations while dyads engaged in free conversations and found that 'ordinary' head motion occured within a range of 1.8 – 3.7 Hz, while 'fast' movements were characterised as above 3.7 Hz. Although their study was based on a very small sample size (*N* = 4), Hadar et al.'s (1983b) account suggests that the decoupling we observed was within a normal to fast range for head movements in conversation. Therefore, given the wealth of evidence that people spontaneously coordinate other movements (e.g. Bernieri, Davis, Rosenthal, & Knee, 1994; Grammer et al., 1998; Ramseyer & Tschacher, 2010; Schmidt et al., 2012), it is surprising we should see active decoupling of head movements at typical frequencies for conversation.

Although we found that the pattern of decoupling was present in all head signals, there was no evidence for the same effect in signals from the torso, which suggests that it is specific to head movements. Our findings are consistent with previous research showing that head and body movements play different roles in interpersonal coordination and show different patterns. For example, Ramseyer & Tschacher (2014) separately examined head and body motion energy from videotapes of psychotherapy sessions. They calculated synchrony in terms of the cross-correlation between patient and psychotherapist, and found that synchronised body motion predicted different outcomes of the therapy session compared to synchronised head motion. Another study by Shockley et al. (2003) used Polhemus motion sensors to capture the head and hip movements of participants engaged in a picture description task. Although they found that participants entrained to one another's bodily sway during the task (assessed through cross-recurrence analysis), they did not find any evidence of entrainment in head displacement data. The authors attributed this null result to disruption in head movements associated with speaking and vocal gesturing. However, we think that the active divergence of head movements in our data could be interpreted as more than just disruption.

We speculate that the active decoupling of head movements we observed in real conversations could be associated with signalling behaviour that accompanies speech. We do not think that the lack of coordination could be due to speech itself, because speech and laughter rhythms occur at around 5 Hz, which is at the very top end of the range where we observed decoupling. Also, if decoupling were due specifically to speech then we would not expect to see the same decoupling during monologue (when only one person was speaking), compared to dialogue (when both people spoke). Instead, we think that the decoupling could be a result of social signalling. Natural head movements at a similar frequency range (2-7 Hz) have been shown to play a direct role in the interpretation of speech (Chandrasekaran et al., 2009; Morrill et al., 2012). Munhall et al. (2004) gave participants a speech-in-noise task and found that they were better at identifying speech when the speech was accompanied by natural head motion from an animated character. Consistent with this data, Hadar et al. (1985) suggested that one person's head movements in a conversation may be signalling their communicative intentions as well as synchronising with the other person. Head signals may include nodding for 'yes' and shaking for 'no', as well as more subtle gestures to indicate impatience, scepticism or interest (Duncan, 1972; Hadar et al., 1983b; Heylen, 2006; McClave, 2000). Hadar et al. (1985) propose that such signalling leads to patterns of dissimilarity in people's head movements. This could be one way to interpret the divergence of head movements in our data at the 1.5 – 5 Hz range, as well as Shockley et al.'s (2003) similar finding that head movements did not coordinate during a picture description task. Therefore, we tentatively suggest that the head decoupling in our data could be associated with dissimilarities arising from the head signals that accompany conversations. However, further research would be needed to test this hypothesis directly.

The decoupling of head movements in our data resonates with other research showing patterns of divergence in social interactions. Firstly, research by Healey et al. (2014) revealed that patterns of syntactic divergence also occur in conversations. Using data from annotated conversations of over 2100 participants, they tested for the repetition of words and syntactic structures between the turns taken in conversation. Real conversations were compared to pseudo conversations, which were created by randomly re-pairing turns from the same conversation. Interestingly, the results showed that people repeated each other's syntactic structures significantly less in real conversation than chance levels would predict, although they may use the same words in a different syntactic structure. The authors suggested that this pattern of divergence could be associated with creating contrasts, evaluations and elaborations that sustain the momentum of the conversation. Thus, people may not always mirror one another in their use of language, but may actively diverge in order to signal particular speech meanings. Another recent study also suggests that partners may show neural decoupling during imitative interactions (Konvalinka et al., 2014). Participants completed a finger-tapping task where they had to synchronise either with another real person, or a computer. Neural activation was measured using EEG, and a computer classifier aimed to distinguish real interactions from computer interactions. For eight out of nine dyads tested, the classifier was able to distinguish real interactions by frontal alpha suppression, which only occurred in one member of each pair. Alpha suppression is associated with planning and control, and so the authors interpret the results as evidence that participants spontaneously adopted distinct leader and follower roles. This is a noteworthy result, because previous research using dual-EEG, dual-fMRI and dual-fNIRS techniques (known as 'hyperscanning') has suggested that when two people engage in a cooperative task together, their brains show similar patterns of activation (Cui, Bryant, & Reiss, 2012; Dumas, Lachat, Martinerie, Nadel, & George, 2011; Konvalinka & Roepstorff, 2012). In contrast, Konvalinka et al's (2014) study suggests that when people aim to coordinate their actions, this may involve different patterns of neural activation in each person in order to achieve synchrony at the behavioural level. Overall, multiple areas of research are beginning to highlight ways in which behaviour systematically diverges during conversations. Such divergence could therefore be important to incorporate into theoretical models of interpersonal coordination. In particular, our results suggest that mimicry of head movements in conversations may be accompanied by active dissimilarity at movement frequencies between 1.5 and 5 Hz.

5.6.2 Coordination of Head and Body Movements at Lower Frequencies

As well as active decoupling of head movements, our results also showed significant coupling of head and body movements at frequencies below 1.5 Hz. These results are consistent with Fujiwara and Daibo's (2016) earlier findings using the pseudo interaction paradigm. The frequency range below 1.5 Hz is roughly consistent with timescales traditionally associated with behavioural mimicry (e.g. Chartrand & Bargh, 1999; Stel et al., 2009). Therefore, we interpret cross-wavelet coherence in this range as evidence for mimicry when looking at the head or torso rotation signals (yaw, pitch and roll). When looking at overall motion energy, we do not know participants' movements were spatially matching, and so cross-wavelet coherence in motion energy is usually interpreted as synchrony (e.g. Fujiwara & Daibo, 2016; Schmidt et al., 2014). By looking at both motion energy and individual rotation signals, we found interesting differences in the way that dyads coordinated their head and body movements at the low frequencies below 1.5 Hz.

Head coordination in this range was mainly characterised by cross-wavelet coherence in head pitch (nodding). Across the different head signals, head pitch showed the greatest coherence relative to chance, and this was significant for all frequencies between 0.07 and 1.1 Hz. In addition, head pitch was the only head signal where the signal power between speakers and listeners was the same. In the other signals, the speaker had greater signal power than the listener, and this difference in the overall amount of movement could have reduced the likelihood of coordination occurring. The fact that speakers and listeners had similar signal power in head pitch suggests that the coordination in this signal was not simply due to movements such as looking downwards at the photograph, since this would occur much more often in the speaker than the listener (who could not see the image). Therefore we suggest that the coordination in the head pitch signal reflected coordinated nodding. Head nodding is known to be an important communicative back-channel for listeners in conversations (Dittmann & Llewellyn, 1969; Hadar et al., 1985; Heylen et al., 2011). Previous research has also suggested that listeners are more likely to nod than make other types of head tilt (Hadar et al., 1983b). Our findings add to this literature by suggesting that participants are more likely than chance to mimic each other's nods during a conversation. Importantly, this pattern did not differ between monologue and dialogue parts of the interaction, which suggests that nodding is not merely a back-channel used by listeners (during monologue) but can also be a mode of coordination between two people engaged in dialogue. In addition to the strong coordination in head pitch, we found that head yaw was significantly more coherent than chance at around 0.1 Hz (a period of 10 seconds). This suggests that coordination of head movements is not restricted to nodding, and that turns of the head may coordinate at a slow timescale. Coordination at 0.1 Hz is at the edge of the expected range for mimicry, and is twice as slow as the rhythm of breathing during conversation (McFarland, 2001). Therefore this peak in coordination is unlikely to be linked to speech, and may possibly reflect a different process to coordination in head pitch.

In contrast to head coordination, torso coordination at frequencies below 1.5 Hz seemed to be characterised by synchronisation in overall motion energy more than one specific rotation channel. This pattern is consistent with the view that behavioural synchrony does not have to involve matching actions (Bernieri & Rosenthal, 1991), and it suggests that natural conversations may involve quite loose matching of body movements unlike the precise mirroring that is generated by virtual mimicry (Chapters 3 & 4). In addition, real interactions showed greater coherence in torso pitch (forward/back lean) compared to pseudo interactions. From the present analyses, it is unclear whether this indicates that both people tended to lean forward (or back) together, or whether they followed a compensatory pattern where one person leans forward when the other leans back (Burgoon, Dillman, & Stem, 1993; Condon & Ogston, 1966). Future analyses could examine the interpersonal distance between participants as a possible indicator of coordination and feelings of interpersonal closeness (Aron et al., 1992; Ashton-James et al., 2007) .

5.6.3 Methodological Implications and Future Directions

Our data also has several methodological implications. Firstly, we found that priming did not have an effect on the levels of coordination in real interactions. In Chapter 1, we suggested that subliminal priming could possibly provide a way to manipulate the levels of mimicry in an interaction without having to explicitly instruct anyone to mimic. This suggestion was based on previous results which show that prosocial and antisocial primes can reliably alter levels of automatic imitation in finger-tapping laboratory tasks. It remained to be tested whether this robust effect would translate into naturalistic situations where the head and body are tracked instead of finger movements. Although our pilot dataset suggested that there might be an effect of priming, we found null results in the final pre-registered analyses which are reported here. This suggests that priming does not have a strong effect on levels of cross-wavelet coherence between two people's head or torso movements, which could mean that priming will not be an effective way to manipulate natural mimicry levels between two naïve participants. However, it is possible that we would have found effects of priming if we had looked at other types of mimicry (e.g. facial mimicry) or used a different measure of interpersonal coordination (e.g. observer ratings of mimicry in video clips).

Secondly, our results provide insight into the natural parameters of interpersonal coordination which could be used to generate more socially realistic virtual characters. This was one of the main motivations for the present study. Our mimicry studies in Chapters 3 and 4 suggested that being precisely mirrored by a virtual character at a fixed time delay cannot generate strong effects on rapport or trust. The findings from the present study could be applied to generate a mimicry algorithm that is closer to the way people naturally mimic each other in conversations. Specifically, our results suggested that mimicry of head movements is naturally restricted to motion frequencies below 1.5 Hz. Whereas our virtual characters previously mimicked the participant's head movements exactly at all motion frequencies, virtual mimicry might be improved in future studies by creating an algorithm that only mimics head movements with frequencies below 1.5 Hz. This could be achieved by blending the relevant frequencies from the participant's motion into a pre-recorded animation that is standard across all participants. By using a pre-recorded animation, this would also create a *lack* of contingency between their head movements at frequencies above 1.5 Hz, where we observed decoupling in real interactions. Importantly for virtual mimicry paradigms, we found that there were virtually no differences in the pattern of interpersonal coordination between monologue and dialogue interactions. This suggests that we can study mimicry in artificial conversations where a participant and a virtual character take turns to speak in monologue without the need to generate realistic dialogue from the virtual character, which is currently very difficult.

Finally, there are many ways that the current dataset or similar data could be explored further. Here, it was not possible to explore all possible analyses and we decided to focus on cross-wavelet coherence of head movements. It would be valuable to examine the phase relationship between participants in more detail (Fujiwara & Daibo, 2016; Issartel et al., 2006; Schmidt et al., 2014). In this study, we did not carry out any tests on the phase relationship, although we visually inspected the phase data for head motion above and below 1.5 Hz (Figures 5-6 to 5-9). Visual inspection did not suggest any consistent leader-follower dynamics existed between speaker and listener. However, more in-depth testing could reveal whether phase-locking is greater than chance at certain phase offsets or certain time-points in the interaction. It would also be possible to test whether levels of coherence change over the time course of a trial or the whole experiment, consistent with earlier microanalysis studies (Condon & Ogston, 1966; Kendon, 1970).

Another avenue would be to examine the relationship between cross-wavelet coherence and ratings of rapport or trust made by the people interacting (Ramseyer & Tschacher, 2011; Ramseyer & Tschacher, 2014). In the study reported here we did not analyse the relationship between cross-wavelet coherence and rapport, trust or other social outcomes traditionally linked to mimicry (which were tested in Chapters 3 and 4). We would predict that cross-wavelet coherence might be positively associated with feelings of rapport and trust, and that coordinated behaviour might lead to increased rapport and trust in a pre-post experiment comparison, although this is a tentative prediction given the mixed evidence that being mimicked by a confederate or virtual character can generate rapport and trust (Chapters 1, 3 and 4). At present, it is also unclear to what extent cross-wavelet coherence can be linked to more traditional measures of social coordination, and so this could help to build up a better understanding of how to interpret cross-wavelet coherence values.

Lastly, we also collected audio data that were not analysed in the present study, which could potentially provide insight into the links between movement coordination and speech content. For example, it could be possible to test whether decoupling in head movements is related to particular prosodic, syntactic or lexical features of speech. Overall, current motion technologies and automatic analysis techniques mean that there could be great value in collecting rich datasets such as those reported here, because they provide many opportunities for investigating interpersonal coordination at a variety of levels. With many new analysis options being explored in current literature, it will be valuable for researchers to share their datasets and use pre-registration platforms to clearly document analyses which are exploratory versus confirmatory.

5.7 Conclusions

Our results from wavelet analysis of head movements provide new insights into the frequency characteristics of interpersonal coordination. We showed that people's head movements in natural conversations involve reliable patterns of dissociation at frequencies between 1.5 and 5 Hz as well as coordination at frequencies below that range. This pattern of decoupling was specific to head movements, and we speculate that it could reflect head signals that accompany speech. We also found that coordination in head movements in conversations are largely characterised by mimicry of nods, whereas coordination in body movement might reflect synchronised posture shifts. The present study builds on a growing literature exploring interpersonal coordination through automatic motion capture and spectrum analysis. Future research could use the insights generated from this approach in order to improve the social realism of virtual characters for future tests of interpersonal coordination.

Chapter 6. Discussion

6.1 Summary of Experimental Chapters

The studies reported in this thesis used novel methods to examine the role of mimicry in rapport and trust. As well as developing new methods, this thesis aimed to rigorously test the claim that being mimicked leads to rapport and trust towards the mimicker. The results from our studies do not support a strong version of this claim, and suggest that mimicry of head and body movements alone has small effects on positive social outcomes including trust and rapport.

In Chapter 2, we developed a new behavioural task for measuring implicit trust towards a specific person. Existing questionnaire measures and economic games have limitations for measuring trust towards a specific stranger in an experimental setting. The virtual maze task measures implicit trust behaviour towards a target person in terms of approaching them for advice and choosing to follow their advice. Across two experiments in VR we showed that the virtual maze task was sensitive to differences between trustworthy and untrustworthy characters, and yielded larger effects than an economic investment game. The task was also sensitive to specific trust towards each character, whereas the investment game partly reflected individual differences in trusting people more generally. In a third experiment, we demonstrated how the task could be adapted for traditional desktop displays, although this format yielded smaller effects.

Chapters 3 and 4 investigated the social effects of being mimicked using a novel virtual mimicry paradigm. Participants took turns to describe photographs with a virtual character that either mirrored their head and body motion with a set delay, or did not mimic. This paradigm provides a very strict test of the effects of being mimicked, while keeping all other social behaviours constant. In Chapter 3, we reported a relatively exploratory study which tested the effect of being mimicked on a range of social outcomes including rapport, trust, similarity, smoothness of the interaction, and self-other overlap. We also manipulated the time delay in mimicry. Participants rated significantly greater rapport towards a virtual character who mimicked them, compared to a virtual character who made pre-recorded movements. However, this effect was not modulated by the timing of mimicry and we found no other significant effects of being mimicked.

In Chapter 4, we aimed to see if the significant effect of virtual mimicry on rapport could be replicated using a more rigorous pre-registered design. In this study, we also compared the effects of virtual mimicry across in-group and out-group interactions. We did not find any significant effects of being mimicked on ratings of rapport or trust, or implicit behaviour in the virtual maze task. There were also no significant interactions between mimicry and group membership. Overall, our results across the two virtual mimicry studies suggest that strict mimicry of head and body movements alone does not have strong effects on rapport and trust. This could be because virtual mimicry is currently unable to replicate the spontaneous patterns of mimicry that may occur in natural social interactions.

Chapter 5 moved away from a strict emphasis on mimicry in order to investigate the natural parameters of interpersonal coordination more broadly. We used high-resolution motion capture to record the head and body movements of dyads while they took turns to describe photos to each other. Following a pre-registered analysis pathway, we used wavelet analysis to examine levels of interpersonal coordination in different head and body signals across a spectrum of motion frequencies. Participants showed coordination in head nodding and body motion energy at frequencies traditionally associated with mimicry. However, at motion frequencies above 1.5 Hz, participants coordinated their head movements significantly *less* than expected from chance levels. This suggests that natural conversations may involve systematic decoupling of head movements alongside spontaneous mimicry.

In this chapter, I will first discuss the implications of our results for current theories of mimicry. In the following sections, I will go on to suggest possible neurocognitive models of being mimicked and discuss methodological directions for neuroimaging and virtual reality which could advance future mimicry research.

6.2 Theoretical Implications and Emerging Questions

In Chapter 1, I introduced three main groups of theories about of the social purpose of mimicry. The first set of theories suggested that mimicry is an innate behaviour that has evolved because it is socially adaptive (Meltzoff, 2007a; Meltzoff & Moore, 1997; Pentland, 2010). These theories have received limited empirical support, and are directly challenged by evidence that mimicry behaviour is malleable across different social contexts (Lakin et al., 2003; Wang & Hamilton, 2012) and in response to sensorimotor training. Therefore, we did not aim to test the view that mimicry is innately adaptive. The second set of theories proposes that mimicry is a strategic communication tool. The social glue theory (Dijksterhuis, 2005; Lakin et al., 2003), which dominates the mimicry literature, assumes that mimicry originally evolved to facilitate communication but now serves the adaptive function of creating affiliation and rapport through harmonious interactions. According to the social glue theory, people can use mimicry strategically. A neurocognitive account of the strategic control of mimicry is provided by the STORM model (Wang & Hamilton, 2012). However, a third theoretical view opposes the claim that mimicry has a strategic social function. Proponents of this view suggest that mimicry is an evolutionary by-product of domain-general visuomotor associations between perceived and performed actions (Heyes, 2001, 2010; Ray & Heyes, 2011). Therefore, according to the associative sequence learning (ASL) theory, mimicry itself has no special social purpose.

In this thesis, we aimed to test the social glue theory of mimicry by examining whether being mimicked leads to feelings of rapport and trust towards the mimicker. If being mimicked has positive social consequences, this would suggest that mimicry is an effective social strategy and provide evidence for the social glue theory. However, if mimicry itself cannot generate any positive social effects, this would undermine its utility as a social strategy, consistent with the view that mimicry has no special social purpose.

Across two tightly controlled studies, our results favour the null hypothesis that being mimicked does not lead to positive feelings towards the mimicker. In our first virtual mimicry experiment (Chapter 3), we explored a range of positive responses to mimicry that have been reported in the literature, and only found one significant effect of being mimicked on ratings of rapport towards the mimicker. This effect was small, and we were unable to replicate it in our second virtual mimicry study (Chapter 4). Across both studies, Bayesian analyses showed that our data favoured the null hypothesis for almost all effects. Together, the null results from these strictly controlled VR studies suggest that mimicry of head and torso movements alone cannot generate positive feelings of trust and rapport towards the mimicker. This leads us to doubt a strong version of the social glue theory, in which all kinds of mimicry (including mimicry of postures, mannerisms, facial expressions and other behaviours) are all assumed to create positive social consequences. Instead, our results appear to be consistent with predictions from the ASL theory, insofar as our Bayesian results provide some modest evidence that mimicry may not generate social benefits. However, it is important to note we did not directly compare the social glue and ASL theories, and our data in Chapters 3 and 4 do not provide direct evidence about the ASL theory.

It is also important to acknowledge potential limitations of our virtual mimicry studies for drawing wider conclusions about the effects of being mimicked. Firstly, we used the motion from one pilot participant to animate our non-mimicking characters, rather than ‘yoking’ to the previous participant’s motion (Bailenson & Yee, 2005). This ensured that our non-mimicry condition was exactly the same for all participants and avoided any unusual movements or problems with motion capture being carried across participants. However, this has the limitation that our results might reflect specific aspects of the recorded motion, although we subjectively judged the selected motion capture to be typical of our pilot participants. Secondly, our studies and previous virtual mimicry experiments (Bailenson & Yee, 2005; Hasler et al., 2014; Verberne et al., 2013) involved the minimum possible number of virtual characters (one character per within-participants condition). This has advantages for maximising the exposure time that participants spend interacting with each character and matching different characters on appearance and voice characteristics, whilst minimising the possibility participants might conflate different characters. However, this approach limits the generalisability of the results, and thus future virtual mimicry studies could be improved by increasing the number of character stimuli. Thirdly, our virtual mimicry studies and traditional experiments in the field have tested mimicry versus non-mimicry, with various operationalisations of non-mimicry as a control condition (Chapter 1). This allows us to infer whether mimicry significantly differs from non-mimicry, but not how these two conditions differ from not interacting at all. For example, it has been suggested that non-mimicry might have detrimental social effects (van Baaren, Decety, et al., 2009) rather than representing a baseline. Instead, a non-interaction condition might provide a neutral baseline against which we could test the relative effects of mimicry and non-mimicry.

Despite our null results in Chapters 3 and 4, the dyadic data reported in Chapter 5 show that people spontaneously coordinate their head and body movements significantly more than we would expect by chance. This suggests that mimicry does matter in some way to social interactions, which seems to contradict our earlier findings. Whereas our virtual mimicry studies suggest that mimicry of head and body motion doesn't always work as a social strategy, our dyadic datasets suggest that coordination does spontaneously occur between dyads as a function of true social interaction. This tension between our findings makes it difficult to give a definitive answer as to whether mimicry has strategic social benefits. If being mimicked does not reliably lead to feelings of rapport and trust, then this raises the question: why do people mimic others and modulate how much they mimic depending on the social context?

There could be several possible answers to this emerging question. Firstly, it could be that mimicry does increase rapport, trust and other positive evaluations of the mimicker under the right conditions. As we have highlighted earlier, it is possible that our virtual mimicry was too artificial to replicate the effects of human mimicry. Our virtual mimicry involved mirror-matching of head and body motion, with head motion being most salient because the body of the virtual mimickers was mostly obscured from view. This type of mimicry might deviate too far from how people mimic one another in real life. Our results from Chapter 5 indicate that head mimicry may be restricted to certain frequencies and that real interactions also involve active decoupling of head movements at high frequencies. The positive effects of mimicry may depend on conforming to these movement parameters and avoiding mimicry when dissimilarity is more appropriate. In other words, our virtual mimickers may have mimicked the wrong head parameters for achieving positive effects.

The positive effects of mimicry might also depend on copying other actions besides head position or overall posture. In most traditional studies of mimicry, a confederate mimicker copies the actions, gestures and mannerisms of the participant. In particular, face-touching and foot-tapping are highlighted as commonly mimicked actions (Chartrand & Bargh, 1999; Stel, van Baaren, et al., 2010; van Baaren et al., 2003). Therefore, it is possible we would have found positive effects of virtual mimicry if our virtual characters also matched the type of arm or leg movements made by participants. In one existing virtual reality study, this approach had a significant effect on empathy towards the mimicker (Hasler et al., 2014). Therefore, it is possible that we needed to implement different mimicry conditions in order to generate positive social effects of being mimicked. Under this explanation, the social glue theory would need to be elaborated to specify the conditions under which mimicry can lead to positive effects.

Secondly, it could be that mimicry is still a worthwhile social strategy even if it leads to small effects. In our virtual mimicry studies, we strictly tested mimicry of head and torso movements in isolation. Previous research using confederates has also aimed to manipulate behavioural mimicry while keeping other social signals constant (e.g. Chartrand & Bargh, 1999; Redeker, Stel, & Mastop, 2011; van Baaren, Holland, Kawakami, & van Knippenberg, 2004). However, we know that mimicry in real life may be accompanied by many other forms of coordinated behaviour and social cues such as smiling and making eye contact (Bernieri & Rosenthal, 1991; Pentland, 2010; Wallbott, 1995; Wang, Newport, et al., 2011). It is possible that all of these social cues have an additive effect and increase the likelihood of affiliation and rapport. Therefore, even if mimicry itself has small and inconsistent effects, it could still be strategically beneficial to mimic others. This could be especially true considering evidence that mimicry is an unconscious and automatic tendency, rather than a wilful exertion (Brass et al., 2009; Chartrand & Lakin, 2013).

Both of the accounts given so far are consistent with the social glue claim that mimicry can be used as a social strategy, albeit with some caveats. However, a third explanation could be that mimicry is a strategy for understanding and learning from other people's minds, rather than a strategy for increasing rapport. Instead of interpreting the top-down control of mimicry in terms of *enhancing* mimicry towards people with whom we want to affiliate, we could see it in terms of *reducing* mimicry inhibition. Some theorists have suggested that mimicry is 'default' social behaviour (van Baaren, Decety, et al., 2009), and this is consistent with evidence from automatic imitation paradigms showing that we automatically imitate others and have to inhibit this tendency when it conflicts with our goals (e.g. responding to a cue on a computer screen) (Brass, Derrfuss, & von Cramon, 2005; Brass et al., 2009; Spengler et al., 2010). In line with this evidence, some theorists suggest that mimicry may be useful as a 'fast and dirty' default strategy for social learning from others (Whiten, Horner, & Marshall-Pescini, 2005; Whiten et al., 2009). Others have suggested that imitation and other forms of bodily alignment may be useful for understanding other people. Garrod and Pickering (Garrod & Pickering, 2004; Pickering & Garrod, 2013) suggest that bodily alignment contributes to mutual understanding when communicating, while simulation theorists (e.g. Gallese & Goldman, 1998; Oberman & Ramachandran, 2007; Ravenscroft, 1998) have argued that imitating someone makes it possible to simulate their mental state and have a theory of mind (although this is highly contentious; Davies & Stone, 2001; Saxe, 2009; Spaulding, 2012; Stich & Nichols, 1992).

According to these views, mimicry is strategic for social understanding: it benefits the mimicker to imitate others in order to communicate and understand their mental state, but it is not important whether this makes them like the mimicker or not. We could speculate that it may particularly benefit the mimicker learn from and understand the mental states of other ingroup members or socially desirable people. Therefore, in these contexts there may be little need to inhibit mimicry, whereas people might be more inclined to inhibit mimicry and focus on our own goals when interacting with outgroup members or people of lower status. Thus, one explanation for the top-down control of mimicry could be that it is a strategy for social understanding rather than social bonding.

To summarise, the chapters in this thesis give a somewhat contradictory view of mimicry. If being mimicked does not increase rapport or trust (Chapters 3 and 4) then why do people spontaneously coordinate their movements (Chapter 5) and appear to modulate mimicry levels strategically (Lakin & Chartrand, 2003; Wang & Hamilton, 2012)? We cannot give a definitive answer to this emerging question based on the data in this thesis. Our results are most resonant with the view that mimicry may only have positive social effects under certain social conditions. In particular, it may be important to mimic head movements within an appropriate frequency range (below 1.5 Hz, Chapter 5) or to mimic other actions as well. If this account is true, then the social glue theory of mimicry could be elaborated to specify the conditions needed for being mimicked to have positive effects. However, we cannot rule out other possible explanations. It could be mimicking other people has very small positive effects which may add to other social cues that increase rapport and trust. This would imply that social glue effects exist but may be weak. Alternatively, mimicry may be a strategy for social understanding rather than social bonding, and could benefit the mimicker by increasing their ability to understand communicative signals and mental states.

To distinguish between these and other possible answers, we would benefit from a better understanding of the cognitive processes going on in the mimickee. For example, is the unconscious perception of mimicry tuned to particular frequencies or types of behaviour? Is being mimicked a rewarding signal to the mimickee? It may be difficult to answer these kind of questions on the basis of behavioural data alone.

Therefore, in the next section I turn to neural evidence about cognitive processes associated with being mimicked.

6.3 Neurocognitive Models of Being Mimicked

A major direction for future theories of mimicry will be to build more detailed models of cognitive systems involved in being mimicked. At present, we have little insight into the unconscious perception of being mimicked and how this affects the mimickee at a cognitive level. These processes may be subtle and hard to probe with behavioural methods, especially when the mimickee is unaware that they are being mimicked. In this context, neuroimaging data could be invaluable for building cognitive models of being mimicked. For example, we already have detailed neural evidence about the *production* of mimicry, which engages inferior parietal cortex and premotor cortex (Grèzes & Decety, 2001; Iacoboni et al., 1999; Molenberghs, Cunnington, & Mattingley, 2009), commonly referred to as the mirror neuron system (MNS). There are also several detailed neurocognitive models describing how the implementation of mimicry by mirror regions may be subject to top-down social control by medial prefrontal and tempero-parietal regions linked to mentalising (Brass et al., 2009; Cross et al., 2013; Spengler et al., 2010; Stel et al., 2016; Wang & Hamilton, 2012).

Therefore, to address the emerging questions from the previous section, a next step will be to examine current neuroimaging data about being mimicked. As we did not carry out any neuroimaging studies in this thesis, I will draw from the small body of existing research in this area. Firstly I will review neuroimaging studies in which participants were mimicked, imitated or acted in synchrony in order to identify potential neural correlates of being mimicked. I will then draw from existing theoretical ideas to develop three possible neurocognitive models of being mimicked. Each model outlines a possible pathway from mimicry perception to reward activation, and I will discuss their implications for the speculative accounts of mimicry which were outlined in the previous section.

6.3.1 Neural Correlates of Being Mimicked

Owing to the difficulty of studying spontaneous behavioural mimicry under controlled conditions, there is little data on the neural correlates of being mimicked (Guionnet et al., 2012). Only one study has measured a mimickee's neural response to mimicry of their postures and body movements (Hogeveen et al., 2014). However, several other research groups have measured neural activation in response to closely related experiences, including being overtly imitated by a live experimenter or a video stimulus (Brass et al., 2009; Decety, Chaminade, Grèzes, & Meltzoff, 2002; Guionnet et al., 2012), passively observing a mimicry interaction from the perspective of the mimickee (Kühn et al., 2010), and interactional synchrony driven by another person (Cacioppo et al., 2014; Kokal et al., 2011). Here I review available data from these different paradigms in order to identify possible neural systems which may be involved in responding to mimicry. The data (summarised in Table 3) highlights three systems involved in responding to mimicry: (1) a perception-action matching system which recognises when we are being mimicked, (2) a self-other system which relates actions made by self and other, and (3) a reward system associated with positive affect and prosocial behaviour (Figure 6-1). Based on these candidate systems, I will go on to outline three possible neurocognitive models of being mimicked.

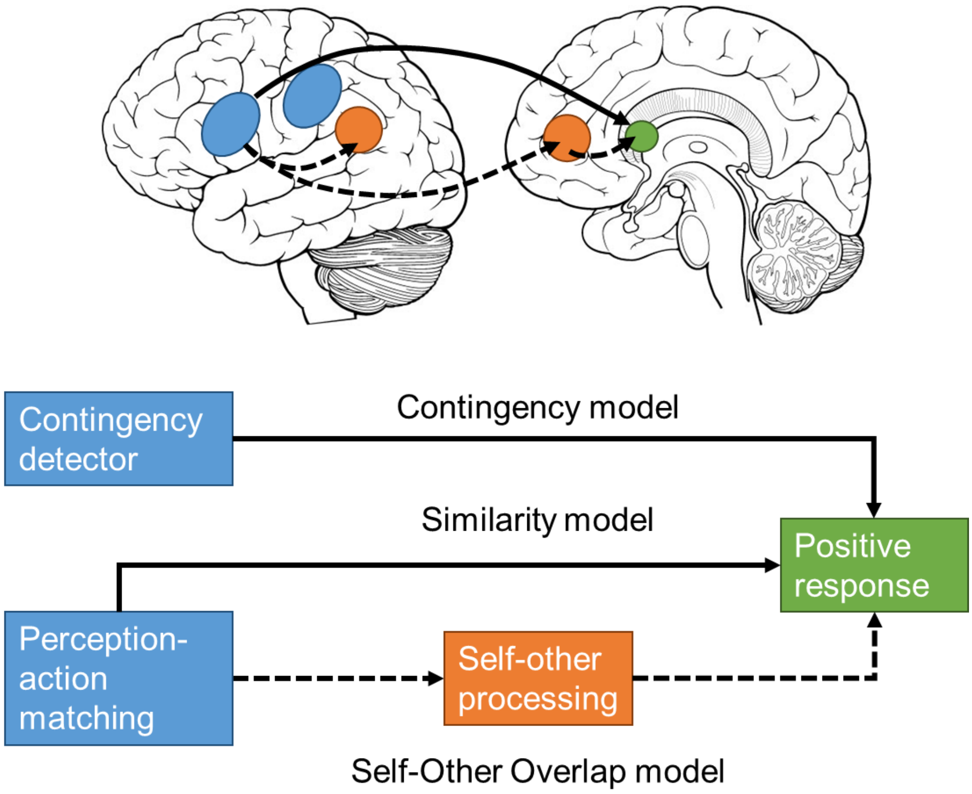

Figure 6-1. Brain regions associated with being mimicked. Unconscious recognition of a perception-action match during mimicry may be associated with MNS activity. Being mimicked increases self-other processing, which may be linked to activity in TPJ and right inferior parietal cortex. Being mimicked is also associated with increased functional connectivity between vmPFC and striatum/insula. Increased activity in striatum and insula may reflect reward and positive responses to being mimicked. The Self-Other Overlap, Contingency and Similarity models predict different cognitive pathways connecting these brain regions.

Perception-action matching. There is extensive evidence that the production of mimicry relies on the mirror system regions of inferior parietal and inferior frontal cortex (Iacoboni et al., 1999; Rizzolatti & Craighero, 2004). These are robustly activated when people produce actions, observe actions and imitate actions (Caspers, Zilles, Laird, & Eickhoff, 2010; Molenberghs et al., 2009). These same regions are also likely to have a role in detecting when someone else is mimicking, because they have the capacity to match observed to performed actions. One study tested this using EEG recordings of the mu-rhythm, a possible marker of MNS function. Hogeveen et al. (2014) took EEG recordings before and after participants completed a rating task. The task involved one of three conditions: social interaction with a mimicking confederate, social interaction with an anti-mimicking confederate (where the confederate changed their behaviour if the participant initiated any mimicry), or interaction with a computer. During EEG recording, participants observed video actions. Their mu-rhythm suppression, which is thought to reflect activation of the sensorimotor cortex, was measured as an indirect index of MNS activity. The results showed enhanced mu-suppression from pre- to post-test in the mimicry condition. The same increase was not found in the anti-mimicry condition, and the increase was significant relative to the computer condition. These findings suggest that being mimicked during naturalistic social interaction leads to an increase in MNS activity which can be detected during subsequent action observation.

Two neuroimaging studies provide evidence that being imitated leads to activation in the left inferior parietal cortex, a classic region of the MNS (Molenberghs et al., 2009; Rizzolatti & Craighero, 2004). Decety et al. (2002) used PET to measure participants' brain activity in response to deliberately imitating or being imitated by an experimenter. The experimenter and participant each had a set of three small objects to manipulate with their right hand and they could see each other's hands via live video links. In this paradigm, participants knew in advance whether they were about to be imitated or not in each block of the experiment. There was an increase in activity in the left inferior parietal cortex when participants were imitated by the experimenter as well as when they did the imitating. Similar activity was found in recent fMRI study of participants who experienced another person not in their view (actually a computer algorithm) synchronising with them on a computer screen while the participant simply tapped a button (Cacioppo et al., 2014). Compared to experiencing asynchrony, while participants experienced synchrony they showed greater activity in the left inferior parietal cortex. Therefore, converging evidence from mimicry, imitation and synchrony paradigms suggests the MNS is involved in the unconscious recognition of mimicry through perception-action matching.

Relation between self and other actions. Being mimicked also appears to activate several regions associated with self-other processing. Decety et al. (2002) found that being imitated was associated with stronger activation in the right inferior parietal cortex, compared to imitating someone. This region is thought to have a role in self-other discrimination and sensing agency (Decety & Sommerville, 2003; Farrer & Frith, 2002; Ruby & Decety, 2001; Uddin, Molnar-Szakacs, Zaidel, & Iacoboni, 2006). Consistent with this finding, Brass et al. (2009) found significant activation in the TPJ in response to being imitated. In their fMRI study, participants made index or middle finger movements that were congruent or incongruent with a stimulus movement, and either saw the stimulus movement before or after they responded. Similar levels of TPJ activity were observed when the participant was imitated and when they experienced an incongruent stimulus. This pattern of results is consistent with the interpretation that TPJ responds when observed movements are delayed or dissimilar performed movements, suggesting this region is involved in distinguishing between self and other actions or perspectives (Brass et al., 2009; Decety & Sommerville, 2003; Ruby & Decety, 2001; Spengler et al., 2010).

However, other results suggest that being mimicked is associated with increased self-other overlap in frontal regions. Kuhn et al. (2010) set out to investigate the neural correlates of positive responses to mimicry. Specifically, participants in an fMRI scanner passively observed videos of social interactions where they took the first-person perspective of an actor being mimicked or anti-mimicked. Compared to anti-mimicry, mimicry led to increased activity in the mOFC/vmPFC, which correlated with ratings of interpersonal closeness. Therefore, being mimicked may be associated with processes of self-other overlap in mOFC/vmPFC in addition to processes of self-other distinction in TPJ and inferior parietal cortex.

Positive responses to mimicry. Neuroimaging data also highlight a system of reward activation in response to being mimicked. In the study described above, Kuhn et al. (2010) also demonstrated activation in brain areas associated with emotion and reward processing. The mimicry condition was associated with increased functional connectivity between vmPFC and the striatum and mid-posterior insula, regions which are related to positive affective states and emotional salience (Craig, 2005; Kühn et al., 2010; Uddin, 2015). In a different paradigm, Guionnet et al. (2012) used live video links to study neural activity while being imitated in an fMRI scanner. Participants either moved their hands and were imitated by an experimenter, or imitated the experimenter's hand movements. Consistent with the functional connectivity reported by Kuhn et al. (2010), there was greater activation in the left anterior insula when participants were imitated. These findings indicate that a reward network involving the striatum and insula may be activated in connection to vmPFC in response to being mimicked.

Further evidence for the same reward system comes from an fMRI study of synchronous behaviour. Kokal et al. (2011) examined activity in the caudate during a drumming task in which participants experienced a partner drumming in synchrony or asynchrony with them. They found that that ease of drumming was associated with activation in the caudate, a region also active in processing monetary reward. Importantly, caudate activation while drumming in synchrony predicted prosocial behaviour towards the drumming partner at the end of the experiment. These findings provide evidence for a neural link from synchrony-related reward processing to downstream prosocial behaviour, which has previously been found to follow synchronised behaviour and mimicry (Valdesolo & DeSteno, 2011; van Baaren et al., 2004; Wiltermuth & Heath, 2009).

6.3.2 Neurocognitive Models

From initial neuroimaging results, it seems that being mimicked may activate three neural systems, one which detects mimicry (MNS), one which relates self and other actions (TPJ and vmPFC), and one which reflects the positive consequences of mimicry (striatum and insula). However, the small number of data points here makes it hard to develop a cognitive model of how these systems might operate together when someone is being mimicked. To advance the field, we can also draw on our extensive knowledge of brain systems engaged in relevant cognitive processes, in particular perception-action matching, social reward processing and perspective-taking. Numerous studies have shown that imitating other people's actions and observing action engages the MNS (Caspers et al., 2010; Molenberghs et al., 2009, 2009). There is also a large body of literature showing that socially rewarding activities engage the insula, ventral striatum and OFC (e.g. Aharon et al., 2001; Bhanji & Delgado, 2014; Fliessbach et al., 2007; Izuma, Saito, & Sadato, 2008; O'Doherty et al., 2003; Sanfey, Rilling, Aronson, Nystrom, & Cohen, 2003). In addition, several lines of evidence suggest that perspective-taking and other forms of self-other processing engage mPFC and TPJ (Brass et al., 2005; David et al., 2006; Denny, Kober, Wager, & Ochsner, 2012; Heatherton et al., 2006; Ruby & Decety, 2001). Drawing on this knowledge in conjunction with the specific studies of being mimicked and synchronising (reviewed above), we can then begin to develop cognitive models which link together these systems and behavioural evidence from mimickees. Next I outline three possible models which build on existing theoretical ideas in the literature and develop them in relation to the available neural evidence. I will evaluate each model in terms of how well it can account for existing behavioural evidence about how people respond to being mimicked, including new insights from the experimental chapters of this thesis.

Self-Other Overlap model. One possible model linking neural and behavioural responses to being mimicked could depend upon self-other processing. During mimicry, the boundary between self and other is thought to become blurred (Georgieff & Jeannerod, 1998). Ashton-James et al. (2007) have proposed that an increase in self-other overlap mediates the prosocial consequences of being mimicked. The Self-Other Overlap model builds on this cognitive pathway by speculating that when perception-action matching occurs in the MNS, regions involved in self-other processing are activated. In turn, frontal regions associated with interpersonal closeness may activate a reward system involving the insula and caudate, which may lead to an increase in prosocial behaviour (Kokal et al., 2011). Other positive responses to mimicry may also result from this cognitive pathway, although only prosocial behaviour has been previously tested (Ashton-James et al., 2007)

Importantly, the Self-Other Overlap model assumes that being mimicked leads to a general tendency to see oneself as closer to others (Ashton-James et al., 2007), despite neural activation in TPJ and inferior parietal cortex associated with self-other distinction (Brass et al., 2009; J. Decety et al., 2002). Several lines of research suggest that the ability to distinguish self- and other-perspectives is essential for taking another's perspective (Galinsky, Ku, & Wang, 2005; Lamm, Batson, & Decety, 2007), which may be an important process in empathy (Bird & Viding, 2014; Jean Decety & Jackson, 2006) and prosocial cooperation (Galinsky et al., 2005; Maddux et al., 2008). Therefore, this model assumes that mimicry ultimately leads people to see others as more 'like me' (Meltzoff, 2007a, 2007b) and behave more prosocially as a result of this self-other overlap.

The Self-Other Overlap model can account for many of the positive responses to mimicry reviewed in the Chapter 1. In particular, several research groups demonstrated that being mimicked makes people behave prosocially towards others in general, and not just the person mimicking (Ashton-James et al., 2007; Carpenter et al., 2013; Fischer-Lokou et al., 2011; Stel & Harinck, 2011; van Baaren et al., 2004). In fact, no studies have reported social effects of being mimicked which failed to extend to other people beyond the mimicker. Furthermore, being mimicked induces cognitive changes in feelings of interdependence (Redeker et al., 2011; Stel & Harinck, 2011; Stel et al., 2011), social distance (Ashton-James et al., 2007) and convergent thinking (Ashton-James & Chartrand, 2009). These findings are consistent with the suggestion that being mimicked primarily increases self-other overlap, and other consequences are secondary. If people tend to rate mimickers as more likeable, trustworthy or persuasive due to a general prosocial effect rather than a change in their perceptions of the mimicker, this could also explain why mimicry appears to have less robust effects on these ratings compared to prosocial behaviour.

However, Hogeveen et al. (2014) found mimicry did not lead to increased self-other overlap using the IOS scale, which is inconsistent with the model's predictions. We also failed to find any significant effects of virtual mimicry on a similar scale (Chapter 3), suggesting that self-other overlap may be a less robust outcome of mimicry than previously assumed. The Sefl-Other Overlap model also does not explain why the positive effects of mimicry are modulated by mimicker characteristics. If being mimicked primarily increases self-other overlap, it is unclear why participants do not respond positively to mimicry from an outgroup member (Chapter 4; Dalton, Chartrand, & Finkel, 2010; Leander, Chartrand, & Bargh, 2012), higher status person (Dalton et al., 2010) or task-focused individual (Leander et al., 2012). Arguably, increased self-other overlap should have especially notable effects in these interactions, because the initial level of overlap may be lowered.

Contingency model. Whereas the first model proposed that perception-action matching is linked to reward via self-other processing, the Contingency model assumes that detecting contingency between our own actions and the world is intrinsically rewarding and motivating. Under this model, complementary and imitative actions would all be processed in the same way and be equally rewarding. From infancy, the ability to detect contingent caregiver behaviour is found to increase positive affect, self-efficacy and social motivation towards the caregiver (Dunham et al., 1989; Millar, 1988; Watson & Ramey, 1972). The Contingency model therefore proposes that being mimicked leads to positive responses due to the contingency of the mimicker's actions on the mimickee's. This view is supported by a recent study showing that people responded positively to contingent movements regardless of how similar the movements were to their own (Catmur & Heyes, 2013), suggesting that positive responses to mimicry may be attributed to contingency and not behaviour-matching. The MNS may be responsible for detecting this contingency. Several studies provide evidence that mirror associations in the MNS are learned through contingent experience, by demonstrating the MNS can form similar associations between dissimilar actions through repeated contingent experiences (Catmur et al., 2008, 2007; Heyes, 2001). Therefore, when the MNS is active in responding to mimicry it may actually reflect the detection of contingency.

The contingency model would predict that positive affective and social consequences of detecting contingency can be attributed to activation of the neural reward system. However, this system may be tuned to an expected *level* of contingency. Infant studies show that contingent behaviour from a stranger only elicits positive responses when the degree of contingency is similar to their caregiver's behaviour (Bigelow, 1998, 2001). Research in robotics also highlights the importance of 'appropriate' contingency levels in creating realistic social entities (Yamaoka et al., 2006, 2007). Therefore, the Contingency model would also predict that reward is not a fixed response to being mimicked.

In support of the Contingency model, being in synchrony has similar positive effects to being mimicked. In particular, synchronised movement leads to increased liking (Hove & Risen, 2009; Miles et al., 2009) and prosocial behaviour (Reddish et al., 2013; Piercarlo Valdesolo et al., 2010; Wiltermuth & Heath, 2009). Synchrony and mimicry also appear to activate similar reward regions in the brain (Cacioppo et al., 2014; Kokal et al., 2011). Since synchronised movements are characterised by contingency rather than similarity (Catmur & Heyes, 2013), this suggests that contingency may explain these effects of mimicry. The Contingency model is also consistent with the breakdown of positive responses to mimicry in contexts and individuals where a lower level of mimicry is typical (Dalton et al., 2010; Leander et al., 2012).

However, the Contingency model would predict that decreasing the time lag in mimicking should elicit stronger responses by making the contingent nature of the mimicker's actions more salient. Although we and others have shown that shorter time delays make it easier to consciously detect that the mimicker is acting contingently (Chapter 3; Bailenson, Beall, Loomis, Blascovich, & Turk, 2004), we did not find any evidence that shorter time delays are associated with more positive social responses in line with the Contingency model. In addition, it is a matter of debate whether mere contingency is 'enough' or whether the similarity of actions has additional importance. The Contingency model is directly challenged by studies comparing merely contingent behaviour to mimicry: in both infants (Agnetta & Rochat 2004) and adults (Hogeveen et al., 2014; Kulesza et al., 2014), mimicry elicits more positive responses than contingent behaviour or anti-mimicry. People also show a preference for movements that involve the same effector even when there is no temporal contingency (Sparenberg, Topolinski, Springer, & Prinz, 2012), which suggests that similarity of movement may still account for some of the positive effects of being mimicked.

Similarity model. Like the Contingency model, the Similarity model proposes a direct pathway from perception-action matching to reward activation, and makes the claim that the most predictable response from the other person is the one with the highest reward value. In this context, the similarity model assumes that an imitative action is more rewarding than a non-imitative one because the kinematic similarity of imitation makes it easier to predict the imitative pattern of action. This means that imitative actions would be more rewarding than complementary actions. Note that overlearned complementary actions (e.g. the grasp patterns involved in handing a mug to another person) might also be highly predictable and thus rewarding.

There is increasing evidence that the brain is good at prediction in both perception and action (Brown & Brüne, 2012; Bubic, von Cramon, & Schubotz, 2010; Clark, 2013). In line with this evidence, the Similarity model assumes the brain is a predictive system which aims to anticipate future sensory inputs (Friston et al., 2011; Kilner et al., 2007), and which finds predictable inputs rewarding. Within this framework, the MNS is part of a generative model that tries to predict incoming sensory input (Kilner, 2011). Using knowledge of a participant's own action and of the social context, the MNS can generate predictions about what the other person will do and can compare those to the other's actual action. If the other person mimics the participant, the visual input is predictable because it is similar to the participant's own action, leading to a low prediction error signal. However, if the other person does not mimic but instead performs some other action, the visual input is less predictable and the error signal is higher. This means that interacting with someone who mimics leads to less prediction error and more activation of reward-related brain networks, which could induce a positive or prosocial mood.

Like the Contingency model, the Similarity model could also generalise to take into account contextual expectations of mimicry. It has previously been suggested that not being imitated is generally unexpected, and therefore experienced negatively (van Baaren, Decety, et al., 2009). If a participant is in a context where mimicry is likely (e.g. interacting with an in-group member), then their MNS will generate a mimicry prediction and when this matches their visual input, prediction error is low and reward is high. However, if a participant is in a context where mimicry is not likely (e.g. interacting with an outgroup member), then their MNS will predict other actions which are not similar to their own. If the interaction partner does mimic, the visual input concerning their actions will not match the predicted visual input, leading to a high prediction error and low reward. Note that this generalisation would require additional contextual information to modulate what the MNS predicts.

By taking mimicry context into account, the Similarity model is able to explain both positive consequences of being mimicked and the breakdown of these positive consequences in certain contexts. Many of the positive effects of mimicry, such as affiliation (Chartrand & Bargh, 1999; Stel & Vonk, 2010), persuasion (Drury & van Swol, 2005; van Swol, 2003) and perceived smoothness (Chartrand & Bargh, 1999) could be direct consequences of reward activation during social interaction. The suggestion that these positive responses depend on the expectation of being mimicked is also consistent with studies showing that mimicry from an outgroup member, high status person or disaffiliative person challenges our expectations, leading to cognitive resource depletion (Dalton et al., 2010) and negative responses (Leander et al., 2012). The Similarity model is also consistent with data suggesting that individual differences in self-construal mediate whether people respond positively to being mimicked. Considering that self-construal is closely tied to cultural norms (Markus & Kitayama, 1991; Sanchez-Burks et al., 2009), people with strongly independent self-construals may expect to be mimicked less often than people who feel strongly interdependent (Sanchez-Burks et al., 2009;

Stel et al., 2011). Thus, people with independent self-construals may not respond positively to mimicry because they do not predict mimicry will occur.

However, this model is less clear in explaining the link between mimicry and prosocial behaviour. It is unclear why a low prediction error and subsequent reward activation should lead to prosocial responses such as helping other people, and why prosocial behaviour should extend beyond the person mimicking (Van Baaren & Chartrand, 2005; van Baaren et al., 2004). Others have suggested that positive affect may be associated with creative and prosocial cognitive styles (Ashton-James & Chartrand, 2009), but there is no clear evidence for a pathway from reward activation to positive affect to generalised prosocial behaviour. Given that increased prosocial behaviour appears to be one of the more consistent effects of being mimicked, this is a significant limitation of the Similarity model.

6.3.3 Summary

The available neuroimaging and EEG data from mimicry, imitation and synchrony tasks suggested that being mimicked may activate mirror neuron systems, brain regions for self-other processing and reward-related systems. I have outlined three speculative models which link these neural systems to possible cognitive processes that follow being mimicked. The Self-Other Overlap model suggests that recognising a perception-action match in the MNS may lead to neural reward via self-other processing; in contrast, the Contingency model and Similarity models propose a direct link between perception-action matching and reward activation (Figure 6-1). The Contingency model argues that this link depends purely on the temporal contingency of the mimicker's actions on the mimickee's and that the kinematic form of their actions is not relevant. In contrast, the Similarity model suggests that kinematic similarity between mimicker and mimickee movements increases the predictability of the mimicker's behaviour, which reduces prediction error and increases reward. Each model is able to predict some of the reported outcomes of being mimicked. However, none of them fully explains the range of mimicry effects reported in previous literature (Chapter 1). This suggests the effects of being mimicked could be explained by a combination of the models above or other models not outlined here.

However, the models outlined in this section have implications for the different theoretical views presented earlier. In the previous section, the question emerged: if being mimicked does not increase rapport or trust, why do people spontaneously mimic others and modulate mimicry levels strategically? I suggested three possible answers. Firstly, mimicry may have positive social effects under certain conditions. This account would be consistent with the Contingency and Similarity models, which both suggest that the perception of being mimicked by the MNS may be sensitive to particular characteristics of the mimicked action. However, it remains unclear whether contingency or similarity or both are important for the detection of mimicry and downstream positive responses. Secondly, I suggested that mimicking other people may be strategic even if it only has very small positive effects. All of the models in this section propose that being mimicked leads (directly or indirectly) to reward activation, which supports the view that mimicking others has positive effects on them. Thirdly, I suggested that mimicry may be a strategy for social understanding instead of social bonding. This view could fit with the Self-Other Overlap model, which suggests that any social benefits of mimicry are mediated by merging of self and other, similar to the idea that the purpose of mimicry may be to help us to align with and better understand others (Garrod & Pickering, 2004; Pickering & Garrod, 2013). Overall, all three suggestions could be supported by the different neurocognitive models outlined here. Therefore, an important challenge for future research will be to find ways of distinguishing between competing neurocognitive models and theoretical accounts of being mimicked. In the next section I will discuss methodological innovations which could help to drive forward the field.

6.4 Future Methodological Directions

A theme throughout this thesis has been to develop new methods for creating ecologically valid mimicry interactions and measuring the downstream effects of mimicry.

In order to advance the field and overcome the limitations associated with traditional studies of mimicry using confederates, we will need to continue developing and refining new methodologies. In this section I outline two methodological directions that could benefit future research of mimicry and other forms of nonverbal behaviour in social interactions. The first direction would be to investigate neural responses to being mimicked during naturalistic social interactions using complementary fMRI and near-infrared spectroscopy techniques. The second direction would be use a data-driven machine learning approach to generate realistic mimicry behaviour from virtual characters based on automatic extraction of mimicry parameters from real world social interactions. While each of these directions could offer new research opportunities, a major challenge will be to create effective collaborations across disciplines. Finally, I briefly suggest a broader scope for investigating mimicry beyond the dyadic level.

6.4.1 Neuroimaging of Social Interactions

To gain a more accurate understanding of brain regions and cognitive processes involved in being mimicked, it will be necessary to measure neural responses in participants during true mimicry interactions. Currently, a major challenge for any neuroimaging study of mimicry is generating appropriate behaviour under controlled conditions. The participant must generate behaviour which can be mimicked, but they must also not be aware that the mimicry is occurring. However, most neuroimaging modalities require the participant to keep still, which restricts the range of possible movements they can perform. To overcome this challenge, future neuroimaging studies could take two different approaches.

First, virtual mimicry could be combined with fMRI. Crucially, virtual mimicry paradigms involve very precise control of mimicry timing and may therefore provide suitable manipulations for fMRI. In order to translate virtual mimicry into the scanner setting, it would be necessary to use a non-magnetic motion tracking system to record the participant's movements and drive the virtual character's behaviour. Due to the sensitivity of fMRI to motion artifacts, it would also be necessary to restrict the range of head and body movement made by the participant within the scanner. Freedom of movement could be increased by using an optical tracking system to control for motion artifacts (Zaitsev, Dold, Sakas, Hennig, & Speck, 2006), or alternatively hand movements could be the target mimicry (cf. Guionnet et al., 2012). However, the constriction of the fMRI environment might also make it difficult to achieve an ecologically valid social interaction when using virtual characters. Another challenge would be to forge collaborations between researchers with expertise in VR and neuroimaging. Although some studies have already used virtual agents and environments for neuroimaging (e.g. Gould et al., 2007; King, Burgess, Hartley, Vargha-Khadem, & O'Keefe, 2002; Ninaus et al., 2014; Schilbach et al., 2010, 2011; Wilms et al., 2010), these are much rarer than behavioural social psychology studies using VR, where there is greater overlap between disciplines. Therefore, research combining VR and neuroimaging might benefit from more meetings and workshops for collaboration between computer scientists and neuroscientists.

A second option would be to study mimicry in live face-to-face interactions using functional near-infrared spectroscopy (fNIRS). Whereas most neuroimaging modalities require the participant to keep still, which restricts the range of movements that can be mimicked, fNIRS avoids this challenge by measuring haemodynamic responses in the brain using infrared light optodes fitted against the scalp. Since fNIRS is portable and much less sensitive to motion artifacts than fMRI or EEG, participants are able to move freely in a face-to-face interaction. A recent experiment demonstrated this possibility by using fNIRS while participants played the popular dance video game, Dance Dance Revolution (Noah et al., 2015). Participants also completed a version of the game adapted for fMRI, and the researchers confirmed there were equivalent activation patterns between the two methods, consistent with other cross-validations (Irani, Platek, Bunce, Ruocco, & Chute, 2007). However, participants were asked not to touch their face or head while wearing the fNIRS optodes in order to avoid face-touching artifacts; this is a disadvantage for studies of mimicry, as face-touching is a commonly mimicked action (Cheng & Chartrand, 2003; Lakin & Chartrand, 2003; van Baaren et al., 2003). Limited depth of penetration in fNIRS also presents a major challenge for testing the possible role of the neural reward system in neurocognitive models of mimicry, because activity in regions such as the caudate and insula would not be detectable using fNIRS.

Moving forward, an optimal strategy for neuroimaging studies of being mimicked could be to carry out complementary experiments using fMRI and fNIRS (Noah et al., 2015). Whereas fMRI could provide high spatial resolution about brain regions activated by being mimicked, fNIRS provides greater ecological validity to examine mimicry in real-world contexts. Despite this advantage, many studies using fNIRS are still currently restricted to artificial laboratory settings (Piper et al., 2014). However, to use fNIRS in more naturalistic settings would provide valuable neural data that could help to distinguish between possible neurocognitive models for responding to mimicry, such as those outlined in the previous section. Due to the scarcity of neural data from participants being mimicked, these models had to draw from neuroimaging studies which tapped into related processes such as deliberate imitation and behavioural synchrony. In order to generate more detailed and accurate models of the neural and cognitive processes involved in being mimicked, it will therefore be important to exploit neuroimaging methods such as fNIRS to measure responses to being mimicked in naturalistic social settings.

6.4.2 Improving Virtual Mimicry through Machine Learning

Another major opportunity for future research will be to take a data-driven approach to building more socially realistic virtual characters. In our virtual mimicry experiments (Chapters 3 and 4) we have demonstrated the experimental benefits of ‘reverse engineering’ virtual characters in order to generate highly controlled social interactions. However, these studies also highlighted major limitation of the reverse-engineered approach: we currently lack detailed knowledge of many behavioural parameters that might need to be programmed into virtual characters, such as when and how much they should mimic, when it is appropriate to blink or smile, or what combinations of nonverbal behaviours should go together. Until we have a more systematic knowledge of people's natural behaviour it will be very difficult to manually design fully realistic virtual characters that can generate social responses equal (or very close) to human confederates. Virtual agents that do successfully generate feelings such as trust and rapport have been achieved through many iterations of trial-and-error in programming (e.g. Gratch et al., 2006, 2007; Huang, Morency, & Gratch, 2011; Verberne et al., 2013, 2015).

However, there is a current trend in computer science of moving away from trial-and-error to developing more data-driven ways of animating virtual characters. In Chapter 5, we outlined how our data about the frequency parameters of mimicry in naturalistic conversations could be used to inform more naturalistic algorithms for animating virtual characters. We suggested that virtual mimickers might achieve stronger social effects if they only mimicked head movements at frequencies below 1.5 Hz, as this was the boundary we found for mimicry of head movements in natural conversations. Although we think this kind of tweak to our current avatars would be a step in the right direction, to generate a fully realistic avatar this way would rely on further studies of other mimicry parameters and very time-consuming programming to translate observed behaviour into animation. However, these limitations might be overcome by using machine learning to uncover the natural patterns of behaviour in social interactions.

Machine learning is a computing method associated with cognitive simulation, or artificial intelligence (Michalski, Carbonell, & Mitchell, 2013). It involves programming computers with algorithms that can learn from and make predictions about datasets without being explicitly programmed. For example, given a dataset about diseases and their symptoms, a machine learning algorithm could learn to predict the disease a patient has based on their symptoms (Koller & Friedman, 2009). In the context of modelling social interactions, machine learning can extract correlated features from a dataset describing the social interaction and then generate a probabilistic model predicting the behaviour of one person based on the other.

If we wanted to get a model that predicts when and how to mimic a participant, then the first step would be to motion track real people engaging in mimicry during a social interaction and record their data as a time series (as we did in Chapter 5). The next step would be to apply a machine learning algorithm to reduce the dimensionality of the dataset, by 'clustering' the motion data into different possible 'actions' (Gillies, 2009). The final step would be to computationally model the probabilistic relationship between different actions, such as nodding and leaning forwards, or mimicked actions such as one person nodding and the other person nodding at the next time point. Hidden Markov models are suitable for this step, because they are able to model sequential outcomes (such as actions) driven by 'hidden' states within the system (analagous to mental states) (e.g. Mead, Atrash, & Matarić, 2013; Mihoub, Bailly, & Wolf, 2013). The model derived from real world data can then be used to drive the behaviour of a virtual character: if input behaviour from a participant is fed into the model, then the virtual character will be able to respond appropriately to the participant's behaviour in real time. In this way, machine learning may be able to generate virtual character animations that closely approximate real-world behaviour without the need to manually test different behavioural parameters and code them individually (Gillies, 2009).

This kind of data-driven approach to generating social behaviours from virtual characters will require strong interdisciplinary collaboration. Implementing machine learning models and animating virtual characters from scratch is outside of the scope of most social psychology laboratories. On the computing side, theoretical knowledge of social interactions is also needed for the 'human component' of designing and fine-tuning of computer models (Gillies et al., 2016). Both sides may also need to work with technologists and engineers to select and use motion capture equipment suitable for capturing social behaviours. Therefore, for data-driven approaches to drive forward the social realism of avatars as research tools, strong collaborations will be needed across disciplines.

6.4.3 Beyond Dyadic Interactions

Throughout this thesis, we specifically focused on mimicry between adult dyads, reflecting much of the empirical work on mimicry, synchrony and other forms of interpersonal coordination. Here, I would like to briefly draw attention to the broader possibility of studying mimicry at a group level.

As a general rule, experiments and theories in the field of mimicry have focused on dyadic interactions where one person mimics another. Some exceptions include early microanalysis studies of interpersonal coordination in videotapes of many people interacting (e.g. Condon & Ogston, 1966; Kendon, 1970) and more recent motion capture research measuring how much people mimic one another in small groups of three (Feese et al., 2011, 2012). One study by van Swol (2003) also tested the effect of mimicry in a group discussion between one participant and two confederates, where only one confederate mimicked. In addition, van Baaren and Chatrand (2005) demonstrated that prosocial consequences of being mimicked may extend to helping people outside of the dyad. However, very few other studies have examined mimicry production in group settings or the group consequences of being mimicked. A related phenomenon which has received more attention is the contagion of emotions and moods within groups (Chartrand & Lakin, 2013; Neumann & Strack, 2000). Recently, two studies have focused on the production of mimicry towards multiple targets. Cracco et al. (2015) found evidence for the automatic imitation of finger movements two people presented at the same time in a spatial response compatibility (SRC) task. Another study by Tsai, Sebanz and Knoblich (2011) showed that pairs of participants mimic hand actions made by other pairs more than actions made by an individual, leading them to suggest that groups mimic group actions. Together, these finding suggests people may mimic multiple others within a group. This opens up the future possibility of exploring the social consequences of mimicking and being mimicked by multiple people.

6.5 Conclusion

This thesis has used novel methods to rigorously test the claim that being mimicked leads to rapport and trust towards the mimicker, in line with the social glue theory of mimicry. This work has extended the field of mimicry research by introducing new virtual reality paradigms, testing novel factors that could modulate responses to being mimicked, and developing a new wavelet analysis pipeline for examining the rhythmic properties of spontaneous mimicry. In addition we have contributed to improving scientific practice in this field by replicating exploratory studies with strict, pre-registered methods and reporting Bayesian analyses as an alternative to null-hypothesis significance tests. Under strictly controlled conditions in virtual reality, we did not find reliable effects of mimicry on rapport or trust, suggesting that mimicry of head and torso movements alone cannot generate positive social effects as the social glue theory predicts. We found that spontaneous interpersonal coordination in naturalistic conversations involves reliable decoupling in head movements alongside mimicry and synchronisation. Overall, our data suggests that mimicry does not have a straightforward role in creating rapport and trust, and that these outcomes may depend on the specific parameters of mimicry or the presence of other social cues.

References

Abbe, A., & Brandon, S. E. (2013a). Building and maintaining rapport in investigative interviews. *Police Practice and Research*, 1–14. https://doi.org/10.1080/15614263.2013.827835

Abbe, A., & Brandon, S. E. (2013b). The Role of Rapport in Investigative Interviewing: A Review. *Journal of Investigative Psychology and Offender Profiling*, *10*(3), 237–249. https://doi.org/10.1002/jip.1386

Aharon, I., Etcoff, N., Ariely, D., Chabris, C. F., O'Connor, E., & Breiter, H. C. (2001). Beautiful faces have variable reward value: fMRI and behavioral evidence. *Neuron*, *32*(3), 537–551.

Allport, G. W. (1968). *The person in psychology: Selected essays* (Vol. viii). Boston, MA, US: Beacon Press.

Ambady, N., Bernieri, F. J., & Richeson, J. A. (2000). Toward a histology of social behavior: Judgmental accuracy from thin slices of the behavioral stream. In B.-A. in E. S. Psychology (Ed.) (Vol. 32, pp. 201–271). Academic Press. Retrieved from http://www.sciencedirect.com/science/article/pii/S0065260100800064

Ambady, N., & Rosenthal, R. (1992). Thin slices of expressive behavior as predictors of interpersonal consequences: A meta-analysis. *Psychological Bulletin*, *111*(2), 256–274. https://doi.org/10.1037/0033-2909.111.2.256

Ambady, N., & Rosenthal, R. (1993). Half a minute: Predicting teacher evaluations from thin slices of nonverbal behavior and physical attractiveness. *Journal of Personality and Social Psychology*, *64*(3), 431–441. https://doi.org/10.1037/0022-3514.64.3.431

Anderson, D. E., DePaulo, B. M., Ansfield, M. E., Tickle, J. J., & Green, E. (1999). Beliefs About Cues to Deception: Mindless Stereotypes or Untapped Wisdom? *Journal of Nonverbal Behavior*, *23*(1), 67–89. https://doi.org/10.1023/A:1021387326192

Anisfeld, M. (1996). Only Tongue Protrusion Modeling Is Matched by Neonates. *Developmental Review*, *16*(2), 149–161. https://doi.org/10.1006/drev.1996.0006

Arena, M. J., Pentland, S., & Price, D. (2010). Honest Signals - Hard Measures for Social Behavior. *Organization Development Journal*, *28*(3), 11–20.

Armitage, C. J., & Christian, J. (2003). From attitudes to behaviour: Basic and applied research on the theory of planned behaviour. *Current Psychology*, *22*(3), 187–195. https://doi.org/10.1007/s12144-003-1015-5

Aron, A., Aron, E. N., & Smollan, D. (1992). Inclusion of Other in the Self Scale and the structure of interpersonal closeness. *Journal of Personality and Social Psychology*, *63*(4), 596–612. https://doi.org/10.1037/0022-3514.63.4.596

Ashraf, N., Bohnet, I., & Piankov, N. (2006). Decomposing trust and trustworthiness. *Experimental Economics*, *9*(3), 193–208. https://doi.org/10.1007/s10683-006-9122-4

Ashton-James, C. E., & Chartrand, T. L. (2009). Social cues for creativity: The impact of behavioral mimicry on convergent and divergent thinking. *Journal of Experimental Social Psychology*, *45*(4), 1036–1040. https://doi.org/10.1016/j.jesp.2009.04.030

Ashton-James, C. E., van Baaren, R. B., Chartrand, T. L., Decety, J., & Karremans, J. (2007). Mimicry and Me: The Impact of Mimicry on Self–Construal. *Social Cognition*, *25*(4), 518–535. https://doi.org/10.1521/soco.2007.25.4.518

Bachmann, R., & Zaheer, A. (2006). *Handbook of Trust Research*. Edward Elgar Publishing.

Bailenson, J. N., Beall, A. C., Loomis, J., Blascovich, J., & Turk, M. (2004). Transformed Social Interaction: Decoupling Representation from Behavior and Form in Collaborative Virtual Environments. *Presence: Teleoperators and Virtual Environments*, *13*(4), 428–441. https://doi.org/10.1162/1054746041944803

Bailenson, J. N., Blascovich, J., Beall, A. C., & Loomis, J. M. (2001). Equilibrium Theory Revisited: Mutual Gaze and Personal Space in Virtual Environments. *Presence: Teleoperators and Virtual Environments*, *10*(6), 583–598. https://doi.org/10.1162/105474601753272844

Bailenson, J. N., Blascovich, J., Beall, A. C., & Loomis, J. M. (2003). Interpersonal Distance in Immersive Virtual Environments. *Personality and Social Psychology Bulletin*, *29*(7), 819–833. https://doi.org/10.1177/0146167203029007002

Bailenson, J. N., & Yee, N. (2005). Digital Chameleons Automatic Assimilation of Nonverbal Gestures in Immersive Virtual Environments. *Psychological Science*, *16*(10), 814–819. https://doi.org/10.1111/j.1467-9280.2005.01619.x

Bailenson, J. N., Yee, N., Patel, K., & Beall, A. C. (2008). Detecting digital chameleons. *Computers in Human Behavior*, *24*(1), 66–87. https://doi.org/10.1016/j.chb.2007.01.015

Banakou, D., Groten, R., & Slater, M. (2013). Illusory ownership of a virtual child body causes overestimation of object sizes and implicit attitude changes. *Proceedings of the National Academy of Sciences*, *110*(31), 12846–12851. https://doi.org/10.1073/pnas.1306779110

Bargh, J. A. (2012). Priming Effects Replicate Just Fine, Thanks. Retrieved June 15, 2017, from http://www.psychologytoday.com/blog/the-natural-unconscious/201205/priming-effects-replicate-just-fine-thanks

Bargh, John A., Chen, M., & Burrows, L. (1996). Automaticity of social behavior: Direct effects of trait construct and stereotype activation on action. *Journal of Personality and Social Psychology*, *71*(2), 230–244. https://doi.org/10.1037/0022-3514.71.2.230

Bascandziev, I., & Harris, P. L. (2014). In beauty we trust: Children prefer information from more attractive informants. *British Journal of Developmental Psychology*, *32*(1), 94–99. https://doi.org/10.1111/bjdp.12022

Bavelas, J., & Healing, S. (2013). Reconciling the effects of mutual visibility on gesturing: A review. *Gesture*, *13*(1), 63–92. https://doi.org/10.1075/gest.13.1.03bav

Bavelas, Janet B., Black, A., Lemery, C. R., & Mullett, J. (1987). Motor mimicry as primitive empathy. In N. Eisenberg & J. Strayer (Eds.), *Empathy and its development* (pp. 317–338). New York, NY, US: Cambridge University Press.

Bavelas, J.B., Black, A., Lemery, C. R., & Mullett, J. (1986). “I Show How You Feel”. Motor Mimicry as a Communicative Act. *Journal of Personality and Social Psychology*, *50*(2), 322–329.

Bekkering, H., & Prinz, W. (2002). Goal representations in imitative actions. In K. Dautenhahn & C. L. Nehaniv (Eds.), *Imitation in animals and artifacts* (pp. 555–572). Cambridge, MA, US: MIT Press.

Bekkering, H., Wohlschlager, A., & Gattis, M. (2000). Imitation of Gestures in Children is Goal-directed. *The Quarterly Journal of Experimental Psychology Section A*, *53*(1), 153–164. https://doi.org/10.1080/713755872

Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B. Methodological*, *57*(1), 289–300.

Ben-Ner, A., & Halldorsson, F. (2010). Trusting and trustworthiness: What are they, how to measure them, and what affects them. *Journal of Economic Psychology*, *31*(1), 64–79. https://doi.org/10.1016/j.joep.2009.10.001

Bentivoglio, A. R., Bressman, S. B., Cassetta, E., Carretta, D., Tonali, P., & Albanese, A. (1997). Analysis of blink rate patterns in normal subjects. *Movement Disorders: Official Journal of the Movement Disorder Society*, *12*(6), 1028–1034. https://doi.org/10.1002/mds.870120629

Berg, J., Dickhaut, J., & McCabe, K. (1995). Trust, Reciprocity, and Social History. *Games and Economic Behavior*, *10*(1), 122–142. https://doi.org/10.1006/game.1995.1027

Bernieri, F. J. (1988). Coordinated movement and rapport in teacher-student interactions. *Journal of Nonverbal Behavior*, *12*(2), 120–138. https://doi.org/10.1007/BF00986930

Bernieri, F. J., Davis, J. M., Rosenthal, R., & Knee, C. R. (1994). Interactional Synchrony and Rapport: Measuring Synchrony in Displays Devoid of Sound and Facial Affect. *Personality and Social Psychology Bulletin*, *20*(3), 303–311. https://doi.org/10.1177/0146167294203008

Bernieri, F. J., Gillis, J. S., Davis, J. M., & Grahe, J. E. (1996). Dyad rapport and the accuracy of its judgment across situations: A lens model analysis. *Journal of Personality and Social Psychology*, *71*(1), 110–129. https://doi.org/10.1037/0022-3514.71.1.110

Bernieri, F. J., & Rosenthal, R. (1991). Interpersonal coordination: Behavior matching and interactional synchrony. In R. S. Feldman & B. Rim (Eds.), *Fundamentals of nonverbal behavior* (pp. 401–432). Paris, France: Editions de la Maison des Sciences de l'Homme.

Bernieri, F. J., Steven, J., & Rosenthal, R. (1988). Synchrony, pseudosynchrony, and dissynchrony: Measuring the entrainment process in mother-infant interactions. *Journal of Personality and Social Psychology*, *54*(2), 243–253. https://doi.org/10.1037/0022-3514.54.2.243

Bhanji, J. P., & Delgado, M. R. (2014). The social brain and reward: social information processing in the human striatum. *Wiley Interdisciplinary Reviews: Cognitive Science*, *5*(1), 61–73. https://doi.org/10.1002/wcs.1266

Bigelow, A. E. (1998). Infants' sensitivity to familiar imperfect contigencies in social interaction. *Infant Behavior and Development*, *21*(1), 149–162. https://doi.org/10.1016/S0163-6383(98)90060-1

Bigelow, A. E. (2001). Discovering self through other: Infants' preference for social contingency. *Bulletin of the Menninger Clinic*, *65*(3), 335–346. https://doi.org/10.1521/bumc.65.3.335.19852

Bird, G., & Viding, E. (2014). The self to other model of empathy: Providing a new framework for understanding empathy impairments in psychopathy, autism, and alexithymia. *Neuroscience & Biobehavioral Reviews*, *47*, 520–532. https://doi.org/10.1016/j.neubiorev.2014.09.021

Blascovich, J., Loomis, J., Beall, A. C., Swinth, K. R., Hoyt, C. L., & Bailenson, J. N. (2002). Immersive Virtual Environment Technology as a Methodological Tool for Social Psychology. *Psychological Inquiry*, *13*(2), 103–124. https://doi.org/10.1207/S15327965PLI1302_01

Blomqvist, K. (1997). The many faces of trust. *Scandinavian Journal of Management*, *13*(3), 271–286. https://doi.org/10.1016/S0956-5221(97)84644-1

Bock, J. K. (1986). Syntactic persistence in language production. *Cognitive Psychology*, *18*(3), 355–387. https://doi.org/10.1016/0010-0285(86)90004-6

Bourgeois, P., & Hess, U. (2008). The impact of social context on mimicry. *Biological Psychology*, *77*(3), 343–352. https://doi.org/10.1016/j.biopsycho.2007.11.008

Brass, M., Derrfuss, J., & von Cramon, D. Y. (2005). The inhibition of imitative and overlearned responses: a functional double dissociation. *Neuropsychologia*, *43*(1), 89–98. https://doi.org/10.1016/j.neuropsychologia.2004.06.018

Brass, M., Ruby, P., & Spengler, S. (2009). Inhibition of imitative behaviour and social cognition. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *364*(1528), 2359–2367. https://doi.org/10.1098/rstb.2009.0066

Brown, E. C., & Brüne, M. (2012). The role of prediction in social neuroscience. *Frontiers in Human Neuroscience*, *6*. https://doi.org/10.3389/fnhum.2012.00147

Brown-Schmidt, S., & Tanenhaus, M. K. (2008). Real-Time Investigation of Referential Domains in Unscripted Conversation: A Targeted Language Game Approach. *Cognitive Science*, *32*(4), 643–684. https://doi.org/10.1080/03640210802066816

Bubic, A., von Cramon, D. Y., & Schubotz, R. I. (2010). Prediction, Cognition and the Brain. *Frontiers in Human Neuroscience*, *4*. https://doi.org/10.3389/fnhum.2010.00025

Burgoon, J. K., Dillman, L., & Stem, L. A. (1993). Adaptation in Dyadic Interaction: Defining and Operationalizing Patterns of Reciprocity and Compensation. *Communication Theory*, *3*(4), 295–316. https://doi.org/10.1111/j.1468-2885.1993.tb00076.x

Cacioppo, S., Zhou, H., Monteleone, G., Majka, E. A., Quinn, K. A., Ball, A. B., … Cacioppo, J. T. (2014). You are in sync with me: Neural correlates of interpersonal synchrony with a partner. *Neuroscience*, *277*, 842–858. https://doi.org/10.1016/j.neuroscience.2014.07.051

Cappella, J. N. (1981). Mutual influence in expressive behavior: Adult–adult and infant–adult dyadic interaction. *Psychological Bulletin*, *89*(1), 101–132. https://doi.org/10.1037/0033-2909.89.1.101

Cappella, J. N., & Planalp, S. (1981). Talk and Silence Sequences in Informal Conversations Iii: Interspeaker Influence. *Human Communication Research*, *7*(2), 117–132. https://doi.org/10.1111/j.1468-2958.1981.tb00564.x

Carpenter, M., Uebel, J., & Tomasello, M. (2013). Being Mimicked Increases Prosocial Behavior in 18-Month-Old Infants. *Child Development*, *84*(5), 1511–1518. https://doi.org/10.1111/cdev.12083

Caspers, S., Zilles, K., Laird, A. R., & Eickhoff, S. B. (2010). ALE meta-analysis of action observation and imitation in the human brain. *NeuroImage*, *50*(3), 1148–1167. https://doi.org/10.1016/j.neuroimage.2009.12.112

Catmur, C., Gillmeister, H., Bird, G., Liepelt, R., Brass, M., & Heyes, C. (2008). Through the looking glass: counter-mirror activation following incompatible sensorimotor learning. *The European Journal of Neuroscience*, *28*(6), 1208–1215. https://doi.org/10.1111/j.1460-9568.2008.06419.x

Catmur, C., & Heyes, C. (2013). Is It What You Do, or When You Do It? The Roles of Contingency and Similarity in Pro-Social Effects of Imitation. *Cognitive Science*, *37*(8), 1541–1552. https://doi.org/10.1111/cogs.12071

Catmur, C., Walsh, V., & Heyes, C. (2007). Sensorimotor Learning Configures the Human Mirror System. *Current Biology*, *17*(17), 1527–1531. https://doi.org/10.1016/j.cub.2007.08.006

Cesario, J., Plaks, J. E., & Tory, E. (2006). Automatic social behavior as motivated preparation to interact. *Journal of Personality and Social Psychology*, *90*(6), 893–910. https://doi.org/10.1037/0022-3514.90.6.893

Chan, D. (2009). So why ask me? Are self-report data really that bad? In *Statistical and methodological myths and urban legends: Doctrine, verity and fable in the organizational and social sciences* (pp. 309–336). New York, NY, US: Routledge/Taylor & Francis Group.

Chandrasekaran, C., Trubanova, A., Stillittano, S., Caplier, A., & Ghazanfar, A. A. (2009). The Natural Statistics of Audiovisual Speech. *PLOS Comput Biol*, *5*(7), e1000436. https://doi.org/10.1371/journal.pcbi.1000436

Chartrand, T. L., & Bargh, J. A. (1999). The chameleon effect: The perception–behavior link and social interaction. *Journal of Personality and Social Psychology*, *76*(6), 893–910. https://doi.org/10.1037/0022-3514.76.6.893

Chartrand, T. L., & Dalton, A. N. (2009). Mimicry: Its Ubiquity, Importance, and Functionality. In E. Morsella, J. A. Bargh, & P. M. Gollwitzer (Eds.), *Oxford Handbook of Human Action* (pp. 458–483). Oxford University Press.

Chartrand, T. L., & Lakin, J. L. (2013). The Antecedents and Consequences of Human Behavioral Mimicry. *Annual Review of Psychology*, *64*(1), 285–308. https://doi.org/10.1146/annurev-psych-113011-143754

Chartrand, T. L., Maddux, W. W., & Lakin, J. L. (2005). Beyond the Perception-Behavior Link: The Ubiquitous Utility and Motivational Moderators of

Nonconscious Mimicry. In R. R. Hassin, J. S. Uleman, & J. A. Bargh (Eds.), *The new unconscious* (pp. 334–361). New York, NY, US: Oxford University Press.

Chartrand, T. L., & van Baaren, R. (2009). Chapter 5 Human Mimicry. In Mark P. Zanna (Ed.), *Advances in Experimental Social Psychology* (Vol. Volume 41, pp. 219–274). Academic Press. Retrieved from http://www.sciencedirect.com/science/article/pii/S006526010800405X

Chase, L. J., & Chase, R. B. (1976). A statistical power analysis of applied psychological research. *Journal of Applied Psychology*, *61*(2), 234–237. https://doi.org/10.1037/0021-9010.61.2.234

Cheng, C. M., & Chartrand, T. L. (2003). Self-Monitoring Without Awareness: Using Mimicry as a Nonconscious Affiliation Strategy. *Journal of Personality*, *85*(6), 1170–1179.

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, *36*(03), 181–204. https://doi.org/10.1017/S0140525X12000477

Clément, F., Koenig, M., & Harris, P. (2004). The ontogenesis of trust. *Mind & Language*, *19*(4), 360–379. https://doi.org/10.1111/j.0268-1064.2004.00263.x

Cohen, J. (1962). The statistical power of abnormal-social psychological research: A review. *The Journal of Abnormal and Social Psychology*, *65*(3), 145–153. https://doi.org/10.1037/h0045186

Cohen, J. (1992). Statistical Power Analysis. *Current Directions in Psychological Science*, *1*(3), 98–101.

Condon, W. S., & Ogston, W. D. (1966). Sound film analysis of normal and pathological behavior patterns. *The Journal of Nervous and Mental Disease*, *143*(4), 338–347.

Cook, J., & Bird, G. (2011). Social attitudes differentially modulate imitation in adolescents and adults. *Experimental Brain Research*, *211*(3–4), 601–612. https://doi.org/10.1007/s00221-011-2584-4

Cook, R., Press, C., Dickinson, A., & Heyes, C. (2010). Acquisition of automatic imitation is sensitive to sensorimotor contingency. *Journal of Experimental Psychology: Human Perception and Performance*, *36*(4), 840–852. https://doi.org/10.1037/a0019256

Couch, L. L., & Jones, W. H. (1997). Measuring Levels of Trust. *Journal of Research in Personality*, *31*(3), 319–336. https://doi.org/10.1006/jrpe.1997.2186

Cox, J. C. (2004). How to identify trust and reciprocity. *Games and Economic Behavior*, *46*(2), 260–281. https://doi.org/10.1016/S0899-8256(03)00119-2

Cracco, E., De Coster, L., Andres, M., & Brass, M. (2015). Motor simulation beyond the dyad: Automatic imitation of multiple actors. *Journal of Experimental Psychology: Human Perception and Performance*, *41*(6), 1488–1501. https://doi.org/10.1037/a0039737

Craig, A. D. (Bud). (2005). Forebrain emotional asymmetry: a neuroanatomical basis? *Trends in Cognitive Sciences*, *9*(12), 566–571. https://doi.org/10.1016/j.tics.2005.10.005

Cross, K. A., Torrisi, S., Reynolds Losin, E. A., & Iacoboni, M. (2013). Controlling automatic imitative tendencies: Interactions between mirror neuron and cognitive control systems. *NeuroImage*, *83*, 493–504. https://doi.org/10.1016/j.neuroimage.2013.06.060

Cruz-Neira, C., Sandin, D. J., & DeFanti, T. A. (1993). Surround-screen Projection-based Virtual Reality: The Design and Implementation of the CAVE. In *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques* (pp. 135–142). New York, NY, USA: ACM. https://doi.org/10.1145/166117.166134

Cui, X., Bryant, D. M., & Reiss, A. L. (2012). NIRS-based hyperscanning reveals increased interpersonal coherence in superior frontal cortex during cooperation. *NeuroImage*, *59*(3), 2430–2437. https://doi.org/10.1016/j.neuroimage.2011.09.003

Dabbs Jr., J. M. (1969). Similarity of gestures and interpersonal influence. *Proceedings of the Annual Convention of the American Psychological Association*, *4*(Pt. 1), 337–338.

Dalton, A. N., Chartrand, T. L., & Finkel, E. J. (2010). The schema-driven chameleon: how mimicry affects executive and self-regulatory resources. *Journal of Personality and Social Psychology*, *98*(4), 605–617. https://doi.org/10.1037/a0017629

Darwin, C. (1859). *On the Origin of Species*. London, UK: Murry.

David, N., Bewernick, B., Cohen, M., Newen, A., Lux, S., Fink, G., … Vogeley, K. (2006). Neural Representations of Self versus Other: Visual-Spatial Perspective Taking and Agency in a Virtual Ball-tossing Game. *Journal of Cognitive Neuroscience*, *18*(6), 898–910. https://doi.org/10.1162/jocn.2006.18.6.898

Davies, M., & Stone, T. (2001). Mental Simulation, Tacit Theory, and the Threat of Collapse.

Davis, S., Nesbitt, K., & Nalivaiko, E. (2014). A Systematic Review of Cybersickness. In *Proceedings of the 2014 Conference on Interactive Entertainment* (p. 8:1–8:9). New York, NY, USA: ACM. https://doi.org/10.1145/2677758.2677780

de Melo, C. M., Carnevale, P., & Gratch, J. (2011). The Impact of Emotion Displays in Embodied Agents on Emergence of Cooperation with People. *Presence: Teleoperators and Virtual Environments*, *20*(5), 449–465. https://doi.org/10.1162/PRES_a_00062

de Melo, C. M., Carnevale, P. J., & Gratch, J. (2013). People’s biased decisions to trust and cooperate with agents that express emotions. In *Trust Workshop* at *Autonomous Agents and Multiagent Systems (AAMAS)*.

de Melo, C. M., Gratch, J., Carnevale, P., & Read, S. (2012). Reverse appraisal: The importance of appraisals for the effect of emotion displays on people’s decision making in a social dilemma. In *Proceedings of the 34th Annual Meeting of the Cognitive Science Society 2012*.

de Melo, C. M., Zheng, L., & Gratch, J. (2009). Expression of Moral Emotions in Cooperating Agents. In *Intelligent Virtual Agents* (pp. 301–307). Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-642-04380-2_32

Decety, J., Chaminade, T., Grèzes, J., & Meltzoff, A. N. (2002). A PET Exploration of the Neural Mechanisms Involved in Reciprocal Imitation. *NeuroImage*, *15*(1), 265–272. https://doi.org/10.1006/nimg.2001.0938

Decety, J., & Sommerville, J. A. (2003). Shared representations between self and other: a social cognitive neuroscience view. *Trends in Cognitive Sciences*, *7*, 527–533. https://doi.org/10.1016/j.tics.2003.10.004

Decety, Jean, & Jackson, P. L. (2006). A Social-Neuroscience Perspective on Empathy. *Current Directions in Psychological Science*, *15*(2), 54–58. https://doi.org/10.1111/j.0963-7214.2006.00406.x

Denny, B. T., Kober, H., Wager, T. D., & Ochsner, K. N. (2012). A Meta-analysis of Functional Neuroimaging Studies of Self- and Other Judgments Reveals a Spatial Gradient for Mentalizing in Medial Prefrontal Cortex. *Journal of Cognitive Neuroscience*, *24*(8), 1742–1752. https://doi.org/10.1162/jocn_a_00233

DePaulo, B. M., Lindsay, J. J., Malone, B. E., Muhlenbruck, L., Charlton, K., & Cooper, H. (2003). Cues to deception. *Psychological Bulletin*, *129*(1), 74–118. https://doi.org/10.1037/0033-2909.129.1.74

Dijksterhuis, A. (2005). Why we are social animals: The high road to imitation as social glue. *Perspectives on Imitation: From Neuroscience to Social Science*, *2*, 207–220.

Dimberg, U., Thunberg, M., & Elmehed, K. (2000). Unconscious facial reactions to emotional facial expressions. *Psychological Science*, *11*(1), 86–89.

Dittmann, A. T., & Llewellyn, L. G. (1969). Body movement and speech rhythm in social conversation. *Journal of Personality and Social Psychology*, *11*(2), 98–106. https://doi.org/10.1037/h0027035

Donath, J. (2007). Virtually Trustworthy. *Science*, *317*(5834), 53–54. https://doi.org/10.1126/science.1142770

Doyen, S., Klein, O., Pichon, C.-L., & Cleeremans, A. (2012). Behavioral Priming: It's All in the Mind, but Whose Mind? *PLoS ONE*, *7*(1), e29081. https://doi.org/10.1371/journal.pone.0029081

Drury, M., & van Swol, L. M. (2005). Are people who mimic others perceived as more friendly, likeable, persuasive, and knowledgeable? Retrieved from http://aran.library.nuigalway.ie/xmlui/handle/10379/1398

Dumas, G., Lachat, F., Martinerie, J., Nadel, J., & George, N. (2011). From social behaviour to brain synchronization: Review and perspectives in hyperscanning. *IRBM*, *32*(1), 48–53.

Duncan, S. (1972). Some signals and rules for taking speaking turns in conversations. *Journal of Personality and Social Psychology*, *23*(2), 283–292. https://doi.org/10.1037/h0033031

Dunham, P., Dunham, F., Hurshman, A., & Alexander, T. (1989). Social Contingency Effects on Subsequent Perceptual-Cognitive Tasks in Young Infants. *Child Development*, *60*(6), 1486–1496. https://doi.org/10.2307/1130937

Durlach, N., & Slater, M. (2000). Presence in Shared Virtual Environments and Virtual Togetherness. *Presence: Teleoperators and Virtual Environments*, *9*(2), 214–217. https://doi.org/10.1162/105474600566736

Elsevier. (n.d.). Registered Reports: A step change in scientific publishing. Retrieved November 24, 2016, from https://www.elsevier.com/reviewers-update/story/innovation-in-publishing/registered-reports-a-step-change-in-scientific-publishing

Ewing, L., Caulfield, F., Read, A., & Rhodes, G. (2014). Perceived trustworthiness of faces drives trust behaviour in children. *Developmental Science*, n/a-n/a. https://doi.org/10.1111/desc.12218

Farrer, C., & Frith, C. D. (2002). Experiencing Oneself vs Another Person as Being the Cause of an Action: The Neural Correlates of the Experience of Agency. *NeuroImage*, *15*(3), 596–603. https://doi.org/10.1006/nimg.2001.1009

Fazio, R. H., & Olson, M. A. (2003). Implicit Measures in Social Cognition Research: Their Meaning and Use. *Annual Review of Psychology*, *54*(1), 297–327. https://doi.org/10.1146/annurev.psych.54.101601.145225

Feese, S., Arnrich, B., Tröster, G., Meyer, B., & Jonas, K. (2011). Detecting posture mirroring in social interactions with wearable sensors (pp. 119–120). Retrieved from http://www.scopus.com/inward/record.url?eid=2-s2.0-80051965159&partnerID=40&md5=6f5339739eb32382a4e3ca0bab6f2036

Feese, S., Arnrich, B., Tröster, G., Meyer, B., & Jonas, K. (2012). Quantifying Behavioral Mimicry by Automatic Detection of Nonverbal Cues from Body Motion. In *Privacy, Security, Risk and Trust (PASSAT), 2012 International Conference on and 2012 International Confernece on Social Computing (SocialCom)* (pp. 520–525). https://doi.org/10.1109/SocialCom-PASSAT.2012.48

Ferguson, C. J., & Heene, M. (2012). A Vast Graveyard of Undead Theories Publication Bias and Psychological Science's Aversion to the Null. *Perspectives on Psychological Science*, *7*(6), 555–561. https://doi.org/10.1177/1745691612459059

Fischer-Lokou, J., Martin, A., Guéguen, N., & Lamy, L. (2011). Mimicry and propagation of prosocial behavior in a natural setting. *Psychological Reports*, *108*(2), 599–605.

Fiske, S. T. (2008). *Social Cognition: From Brains to Culture*. Boston: McGraw-Hill Higher Education.

Fiske, S. T., Cuddy, J., Glick, P., & Xu, J. (2002). A model of (often mixed) stereotype content: Competence and warmth respectively follow from perceived status and competition. *Journal of Personality and Social Psychology*, *82*(6), 878–902. https://doi.org/10.1037/0022-3514.82.6.878

Fiske, S. T., Cuddy, A. J. C., & Glick, P. (2007). Universal dimensions of social cognition: warmth and competence. *Trends in Cognitive Sciences*, *11*(2), 77–83. https://doi.org/10.1016/j.tics.2006.11.005

Fiske, S. T., Xu, J., Cuddy, A. C., & Glick, P. (1999). (Dis)respecting versus (Dis)liking: Status and Interdependence Predict Ambivalent Stereotypes of Competence and Warmth. *Journal of Social Issues*, *55*(3), 473–489. https://doi.org/10.1111/0022-4537.00128

Fliessbach, K., Weber, B., Trautner, P., Dohmen, T., Sunde, U., Elger, C. E., & Falk, A. (2007). Social Comparison Affects Reward-Related Brain Activity in the Human Ventral Striatum. *Science*, *318*(5854), 1305–1308. https://doi.org/10.1126/science.1145876

Fox, J., Arena, D., & Bailenson, J. N. (2009). Virtual Reality: A Survival Guide for the Social Scientist. *Journal of Media Psychology: Theories, Methods, and Applications*, *21*(3), 95–113. https://doi.org/10.1027/1864-1105.21.3.95

Francis, G. (2012). Publication bias and the failure of replication in experimental psychology. *Psychonomic Bulletin & Review*, *19*(6), 975–991. https://doi.org/10.3758/s13423-012-0322-y

Franzen, N., Hagenhoff, M., Baer, N., Schmidt, A., Mier, D., Sammer, G., … Lis, S. (2011). Superior “theory of mind” in borderline personality disorder: An analysis of interaction behavior in a virtual trust game. *Psychiatry Research*, *187*(1–2), 224–233. https://doi.org/10.1016/j.psychres.2010.11.012

Freitag, M., & Traunmüller, R. (2009). Spheres of trust: An empirical analysis of the foundations of particularised and generalised trust. *European Journal of Political Research*, *48*(6), 782–803. https://doi.org/10.1111/j.1475-6765.2009.00849.x

Freud, S. (1921). *Group Psychology and the Analysis of the Ego*. Norton.

Friedman, D., Pizarro, R., Or-Berkers, K., Neyret, S., Pan, X., & Slater, M. (2014). A method for generating an illusion of backwards time travel using immersive virtual reality—an exploratory study. *Perception Science*, *5*, 943. https://doi.org/10.3389/fpsyg.2014.00943

Friston, K., Mattout, J., & Kilner, J. (2011). Action understanding and active inference. *Biological Cybernetics*, *104*(1–2), 137–160. https://doi.org/10.1007/s00422-011-0424-z

Fujiwara, K., & Daibo, I. (2016). Evaluating Interpersonal Synchrony: Wavelet Transform Toward an Unstructured Conversation. *Frontiers in Psychology*, *7*. https://doi.org/10.3389/fpsyg.2016.00516

Galinsky, A. D., Ku, G., & Wang, C. S. (2005). Perspective-Taking and Self-Other Overlap: Fostering Social Bonds and Facilitating Social Coordination. *Group Processes & Intergroup Relations*, *8*(2), 109–124. https://doi.org/10.1177/1368430205051060

Gallese, V., & Goldman, A. (1998). Mirror neurons and the simulation theory of mind-reading. *Trends in Cognitive Sciences*, *2*(12), 493–501.

Garau, M., Slater, M., Pertaub, D.-P., & Razzaque, S. (2005). The Responses of People to Virtual Humans in an Immersive Virtual Environment. *Presence: Teleoperators and Virtual Environments*, *14*(1), 104–116. https://doi.org/10.1162/1054746053890242

Garrod, S., & Pickering, M. J. (2004). Why is conversation so easy? *Trends in Cognitive Sciences*, *8*(1), 8–11. https://doi.org/10.1016/j.tics.2003.10.016

Georgieff, N., & Jeannerod, M. (1998). Beyond consciousness of external reality: a “who” system for consciousness of action and self-consciousness. *Consciousness and Cognition*, *7*(3), 465–477. https://doi.org/10.1006/ccog.1998.0367

Giles, H., & Powesland, P. F. (1975). *Speech style and social evaluation* (Vol. viii). Oxford, England: Academic Press.

Gillath, O., McCall, C., Shaver, P. R., & Blascovich, J. (2008). What Can Virtual Reality Teach Us About Prosocial Tendencies in Real and Virtual Environments? *Media Psychology*, *11*(2), 259–282. https://doi.org/10.1080/15213260801906489

Gillies, M. (2009). Learning Finite-State Machine Controllers From Motion Capture Data. *IEEE Transactions on Computational Intelligence and AI in Games*, *1*(1), 63–72. https://doi.org/10.1109/TCIAIG.2009.2019630

Gillies, Marco, Fiebrink, R., Tanaka, A., Garcia, J., Bevilacqua, F., Heloir, A., … Caramiaux, B. (2016). Human-Centred Machine Learning. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems* (pp. 3558–3565). New York, NY, USA: ACM. https://doi.org/10.1145/2851581.2856492

Gillmeister, H., Catmur, C., Liepelt, R., Brass, M., & Heyes, C. (2008). Experience-based priming of body parts: A study of action imitation. *Brain Research*, *1217*, 157–170.

Glaeser, E., Laibson, D. I., Scheinkman, J. A., & Soutter, C. L. (2000). Measuring Trust. *The Quarterly Journal of Economics*, *115*(3), 811–846. https://doi.org/10.1162/003355300554926

Glaeser, E., Laibson, D., Scheinkman, J., & Soutter, C. L. (1999). *What is Social Capital? The Determinants of Trust and Trustworthiness* (Harvard Institute of Economic Research Working Paper). Harvard - Institute of Economic Research. Retrieved from http://econpapers.repec.org/paper/fthharver/1875.htm

Gould, N. F., Holmes, M. K., Fantie, B. D., Luckenbaugh, D. A., Pine, D. S., Gould, T. D., … Zarate, C. A. (2007). Performance on a Virtual Reality Spatial Memory Navigation Task in Depressed Patients. *American Journal of Psychiatry*, *164*(3), 516–519. https://doi.org/10.1176/ajp.2007.164.3.516

Grammer, K., Kruck, K. B., & Magnusson, M. S. (1998). The Courtship Dance: Patterns of Nonverbal Synchronization in Opposite-Sex Encounters. *Journal of Nonverbal Behavior*, *22*(1), 3–29. https://doi.org/10.1023/A:1022986608835

Gratch, J., Okhmatovskaia, A., Lamothe, F., Marsella, S., Morales, M., Werf, R. J. van der, & Morency, L.-P. (2006). Virtual Rapport. In J. Gratch, M. Young, R. Aylett, D. Ballin, & P. Olivier (Eds.), *Intelligent Virtual Agents* (pp. 14–27). Springer Berlin Heidelberg. Retrieved from http://link.springer.com/chapter/10.1007/11821830_2

Gratch, J., Wang, N., Gerten, J., Fast, E., & Duffy, R. (2007). Creating Rapport with Virtual Agents. In C. Pelachaud, J.-C. Martin, E. André, G. Chollet, K. Karpouzis, & D. Pelé (Eds.), *Intelligent Virtual Agents* (pp. 125–138). Springer Berlin Heidelberg. Retrieved from http://link.springer.com/chapter/10.1007/978-3-540-74997-4_12

Grèzes, J., & Decety, J. (2001). Functional anatomy of execution, mental simulation, observation, and verb generation of actions: A meta-analysis. *Human Brain Mapping*, *12*(1), 1–19. https://doi.org/10.1002/1097-0193(200101)12:1<1::AID-HBM10>3.0.CO;2-V

Grinsted, A., Moore, J. C., & Jevrejeva, S. (2004). Application of the cross wavelet transform and wavelet coherence to geophysical time series. *Nonlinear Processes in Geophysics*, *11*(5/6), 561–566.

Guadagno, R. E., Blascovich, J., Bailenson, J. N., & Mccall, C. (2007). Virtual Humans and Persuasion: The Effects of Agency and Behavioral Realism. *Media Psychology*, *10*(1), 1–22. https://doi.org/10.1080/15213260701300865

Guéguen, N. (2013). Mimicry and Honesty: People Give More Honest Responses to Their Mimicker. *International Journal of Psychological Research*, *6*(1), 53–57.

Guéguen, N., Martin, A., & Meineri, S. (2011). Mimicry and Helping Behavior: An Evaluation of Mimicry on Explicit Helping Request. *The Journal of Social Psychology*, *151*(1), 1–4. https://doi.org/10.1080/00224540903366701

Guéguen, N., Martin, A., Meineri, S., & Simon, J. (2013). Using Mimicry to Elicit Answers to Intimate Questions in Survey Research. *Field Methods*, *25*(1), 47–57. https://doi.org/10.1177/1525822X12449710

Guionnet, S., Nadel, J., Bertasi, E., Sperduti, M., Delaveau, P., & Fossati, P. (2012). Reciprocal imitation: Toward a neural basis of social interaction. *Cerebral Cortex*, *22*(4), 971–978.

Gunnthorsdottir, A., McCabe, K., & Smith, V. (2002). Using the Machiavellianism instrument to predict trustworthiness in a bargaining game. *Journal of Economic Psychology*, *23*(1), 49–66. https://doi.org/10.1016/S0167-4870(01)00067-8

Hadar, U., Steiner, T. J., Grant, E. C., & Rose, F. C. (1983a). Head Movement Correlates of Juncture and Stress at Sentence Level. *Language and Speech*, *26*(2), 117–129. https://doi.org/10.1177/002383098302600202

Hadar, U., Steiner, T. J., Grant, E. C., & Rose, F. C. (1983b). Kinematics of head movements accompanying speech during conversation. *Human Movement Science*, *2*(1–2), 35–46. https://doi.org/10.1016/0167-9457(83)90004-0

Hadar, U., Steiner, T. J., & Rose, F. C. (1985). Head movement during listening turns in conversation. *Journal of Nonverbal Behavior*, *9*(4), 214–228. https://doi.org/10.1007/BF00986881

Hale, J., & Hamilton, A. F. de C. (2015). Does group membership matter for how people respond to mimicry? Open Science Framwork. Retrieved from https://osf.io/n4kb7/

Hale, J., & Hamilton, A. F. D. C. (2016a). Get on my wavelength: The effects of prosocial priming on interpersonal coherence measured with high-resolution motion capture. Retrieved from osf.io/nex72

Hale, J., & Hamilton, A. F. de C. (2016b). Cognitive mechanisms for responding to mimicry from others. *Neuroscience & Biobehavioral Reviews*, *63*, 106–123. https://doi.org/10.1016/j.neubiorev.2016.02.006

Hamilton, A. F. de C. (2008). Emulation and mimicry for social interaction: A theoretical approach to imitation in autism. *The Quarterly Journal of Experimental Psychology*, *61*(1), 101–115. https://doi.org/10.1080/17470210701508798

Harris, C. R., Coburn, N., Rohrer, D., & Pashler, H. (2013). Two Failures to Replicate High-Performance-Goal Priming Effects. *PLOS ONE*, *8*(8), e72467. https://doi.org/10.1371/journal.pone.0072467

Harris, P. L., & Corriveau, K. H. (2011). Young children's selective trust in informants. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, *366*(1567), 1179–1187. https://doi.org/10.1098/rstb.2010.0321

Hasler, B. S., Hirschberger, G., Shani-Sherman, T., & Friedman, D. A. (2014). Virtual Peacemakers: Mimicry Increases Empathy in Simulated Contact with Virtual Outgroup Members. *Cyberpsychology, Behavior, and Social Networking*. https://doi.org/10.1089/cyber.2014.0213

Healey, P. G. T., Purver, M., & Howes, C. (2014). Divergence in Dialogue. *PLOS ONE*, *9*(6), e98598. https://doi.org/10.1371/journal.pone.0098598

Heatherton, T. F., Wyland, C. L., Macrae, C. N., Demos, K. E., Denny, B. T., & Kelley, W. M. (2006). Medial prefrontal activity differentiates self from close others. *Social Cognitive and Affective Neuroscience*, *1*(1), 18–25. https://doi.org/10.1093/scan/nsl001

Heerey, E. A., & Crossley, H. M. (2013). Predictive and Reactive Mechanisms in Smile Reciprocity. *Psychological Science*, 0956797612472203. https://doi.org/10.1177/0956797612472203

Heerey, E. A., & Kring, A. M. (2007). Interpersonal Consequences of Social Anxiety. [Miscellaneous Article]. *Journal of Abnormal Psychology February 2007*, *116*(1), 125–134. https://doi.org/10.1037/0021-843X.116.1.125

Herrmann, A., Rossberg, N., Huber, F., Landwehr, J. R., & Henkel, S. (2011). The impact of mimicry on sales – Evidence from field and lab experiments. *Journal* *of Economic Psychology*, *32*(3), 502–514. https://doi.org/10.1016/j.joep.2011.03.017

Hess, U., & Bourgeois, P. (2010). You smile-I smile: Emotion expression in social interaction. *Biological Psychology*, *84*(3), 514–520.

Heyes, C. (2001). Causes and consequences of imitation. *Trends in Cognitive Sciences*, *5*(6), 253–261. https://doi.org/10.1016/S1364-6613(00)01661-2

Heyes, C. (2010). Where do mirror neurons come from? *Neuroscience & Biobehavioral Reviews*, *34*(4), 575–583. https://doi.org/10.1016/j.neubiorev.2009.11.007

Heyes, C. (2011). Automatic imitation. *Psychological Bulletin*, *137*(3), 463–483. https://doi.org/10.1037/a0022288

Heyes, C., Bird, G., Johnson, H., & Haggard, P. (2005). Experience modulates automatic imitation. *Cognitive Brain Research*, *22*(2), 233–240. https://doi.org/10.1016/j.cogbrainres.2004.09.009

Heyes, C. M., & Ray, E. D. (2000). What Is the Significance of Imitation in Animals? In J. S. R. Peter J.B. Slater Charles T.Snowdon and Timothy J.Roper (Ed.), *Advances in the Study of Behavior* (Vol. 29, pp. 215–245). Academic Press. Retrieved from http://www.sciencedirect.com/science/article/pii/S0065345408601060

Heylen, D. (2006). Head gestures, gaze and the principles of conversational structure. *International Journal of Humanoid Robotics*, *03*(03), 241–267. https://doi.org/10.1142/S0219843606000746

Heylen, D., Bevacqua, E., Pelachaud, C., Poggi, I., Gratch, J., & Schröder, M. (2011). Generating Listening Behaviour. In R. Cowie, C. Pelachaud, & P. Petta (Eds.), *Emotion-Oriented Systems* (pp. 321–347). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-15184-2_17

Hogeveen, J., Chartrand, T. L., & Obhi, S. S. (2014). Social Mimicry Enhances Mu-Suppression During Action Observation. *Cerebral Cortex*, bhu016. https://doi.org/10.1093/cercor/bhu016

Hooghe, M., Reeskens, T., Stolle, D., & Trappers, A. (2009). Ethnic Diversity and Generalized Trust in Europe A Cross-National Multilevel Study. *Comparative Political Studies*, *42*(2), 198–223. https://doi.org/10.1177/0010414008325286

Hosman, L. A., & Wright, J. W. (1987). The effects of hedges and hesitations on impression formation in a simulated courtroom context. *Western Journal of Speech Communication*, *51*(2), 173–188. https://doi.org/10.1080/10570318709374263

Hove, M. J., & Risen, J. L. (2009). It's All in the Timing: Interpersonal Synchrony Increases Affiliation. *Social Cognition*, *27*(6), 949–960. https://doi.org/10.1521/soco.2009.27.6.949

Hsee, C. K., Hatfield, E., Carlson, J. G., & Chemtob, C. (1990). The effect of power on susceptibility to emotional contagion. *Cognition & Emotion*, *4*(4), 327–340. https://doi.org/10.1080/02699939008408081

Huang, L., Morency, L.-P., & Gratch, J. (2011). Virtual Rapport 2.0. In H. H. Vilhjálmsson, S. Kopp, S. Marsella, & K. R. Thórisson (Eds.), *Intelligent Virtual Agents* (pp. 68–79). Springer Berlin Heidelberg. Retrieved from http://link.springer.com/chapter/10.1007/978-3-642-23974-8_8

Hull, J. G., Slone, L. B., Meteyer, K. B., & Matthews, A. R. (2002). The nonconsciousness of self-consciousness. *Journal of Personality and Social Psychology*, *83*(2), 406–424. https://doi.org/10.1037/0022-3514.83.2.406

Iacoboni, M., Woods, R. P., Brass, M., Bekkering, H., Mazziotta, J. C., & Rizzolatti, G. (1999). Cortical Mechanisms of Human Imitation. *Science*, *286*(5449), 2526–2528. https://doi.org/10.1126/science.286.5449.2526

Inzlicht, M., Gutsell, J. N., & Legault, L. (2012). Mimicry reduces racial prejudice. *Journal of Experimental Social Psychology*, *48*(1), 361–365. https://doi.org/10.1016/j.jesp.2011.06.007

Irani, F., Platek, S. M., Bunce, S., Ruocco, A. C., & Chute, D. (2007). Functional near infrared spectroscopy (fNIRS): an emerging neuroimaging technology with important applications for the study of brain disorders. *The Clinical Neuropsychologist*, *21*(1), 9–37. https://doi.org/10.1080/13854040600910018

Issartel, J., Bardainne, T., Gaillot, P., & Marin, L. (2015). The relevance of the cross-wavelet transform in the analysis of human interaction – a tutorial. *Quantitative Psychology and Measurement*, *5*, 1566. https://doi.org/10.3389/fpsyg.2014.01566

Issartel, J., Marin, L., Gaillot, P., Bardainne, T., & Cadopi, M. (2006). A Practical Guide to Time—Frequency Analysis in the Study of Human Motor Behavior: The Contribution of Wavelet Transform. *Journal of Motor Behavior*, *38*(2), 139–159. https://doi.org/10.3200/JMBR.38.2.139-159

Izuma, K., Saito, D. N., & Sadato, N. (2008). Processing of Social and Monetary Rewards in the Human Striatum. *Neuron*, *58*(2), 284–294. https://doi.org/10.1016/j.neuron.2008.03.020

Jacob, C., Guéguen, N., Martin, A., & Boulbry, G. (2011). Retail salespeople's mimicry of customers: Effects on consumer behavior. *Journal of Retailing and Consumer Services*, *18*(5), 381–388. https://doi.org/10.1016/j.jretconser.2010.11.006

Jarosz, A., & Wiley, J. (2014). What Are the Odds? A Practical Guide to Computing and Reporting Bayes Factors. *The Journal of Problem Solving*, *7*(1). https://doi.org/10.7771/1932-6246.1167

Johnson, N. D., & Mislin, A. A. (2011). Trust games: A meta-analysis. *Journal of Economic Psychology*, *32*(5), 865–889. https://doi.org/10.1016/j.joep.2011.05.007

Johnson-George, C., & Swap, W. C. (1982). Measurement of specific interpersonal trust: Construction and validation of a scale to assess trust in a specific other. *Journal of Personality and Social Psychology*, *43*(6), 1306–1317. https://doi.org/10.1037/0022-3514.43.6.1306

Johnston, A. M., Mills, C. M., & Landrum, A. R. (2015). How do children weigh competence and benevolence when deciding whom to trust? *Cognition*, *144*, 76–90. https://doi.org/10.1016/j.cognition.2015.07.015

Johnston, L. (2002). Behavioral Mimicry and Stigmatization. *Social Cognition*, *20*(1), 18–35. https://doi.org/10.1521/soco.20.1.18.20944

Jonas, K. J., & Cesario, J. (2015). How can preregistration contribute to research in our field? *Comprehensive Results in Social Psychology*, 1–7. https://doi.org/10.1080/23743603.2015.1070611

Jones, S. S. (2009). The development of imitation in infancy. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *364*(1528), 2325–2335. https://doi.org/10.1098/rstb.2009.0045

Jose, A. D., & Collison, D. (1970). The normal range and determinants of the intrinsic heart rate in man. *Cardiovascular Research*, *4*(2), 160–167. https://doi.org/10.1093/cvr/4.2.160

Karlan, D. S. (2005). Using Experimental Economics to Measure Social Capital and Predict Financial Decisions. *The American Economic Review*, *95*(5), 1688–1699.

Karremans, J. C., & Verwijmeren, T. (2008). Mimicking attractive opposite-sex others: The role of romantic relationship status. *Personality and Social Psychology Bulletin*, *34*(7), 939–950.

Kendon, A. (1970). Movement coordination in social interaction: Some examples described. *Acta Psychologica*, *32*, 101–125. https://doi.org/10.1016/0001-6918(70)90094-6

Khooshabeh, P., Gandhe, S., McCall, C., Gratch, J., Blascovich, J., & Traum, D. (2011). The Effects of Virtual Agent Humor and Gaze Behavior on Human-Virtual Agent Proxemics. In H. H. Vilhjálmsson, S. Kopp, S. Marsella, & K. R. Thórisson (Eds.), *Intelligent Virtual Agents* (pp. 458–459). Springer Berlin

Heidelberg. Retrieved from http://link.springer.com/chapter/10.1007/978-3-642-23974-8_61

Kilner, J. M. (2011). More than one pathway to action understanding. *Trends in Cognitive Sciences*, *15*(8), 352–357. https://doi.org/10.1016/j.tics.2011.06.005

Kilner, J. M., Friston, K. J., & Frith, C. D. (2007). Predictive coding: an account of the mirror neuron system. *Cognitive Processing*, *8*(3), 159–166. https://doi.org/10.1007/s10339-007-0170-2

King, J. A., Burgess, N., Hartley, T., Vargha-Khadem, F., & O'Keefe, J. (2002). Human hippocampus and viewpoint dependence in spatial memory. *Hippocampus*, *12*(6), 811–820. https://doi.org/10.1002/hipo.10070

King-Casas, B., Tomlin, D., Anen, C., Camerer, C. F., Quartz, S. R., & Montague, P. R. (2005). Getting to Know You: Reputation and Trust in a Two-Person Economic Exchange. *Science*, *308*(5718), 78–83. https://doi.org/10.1126/science.1108062

Kinsbourne, M., & Helt, M. (n.d.). Entrainment, mimicry and interpersonal synchrony. In *Fein, D. (Ed.)The Neuropsychology of Autism*. OUP USA.

Kinzler, K. D., Corriveau, K. H., & Harris, P. L. (2011). Children's selective trust in native-accented speakers. *Developmental Science*, *14*(1), 106–111. https://doi.org/10.1111/j.1467-7687.2010.00965.x

Klein, O., Doyen, S., Leys, C., Gama, P. A. M. de S. da, Miller, S., Questienne, L., & Cleeremans, A. (2012). Low Hopes, High Expectations Expectancy Effects and the Replicability of Behavioral Experiments. *Perspectives on Psychological Science*, *7*(6), 572–584. https://doi.org/10.1177/1745691612463704

Knoblich, G., & Sebanz, N. (2008). Evolving intentions for social interaction: from entrainment to joint action. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, *363*(1499), 2021–2031. https://doi.org/10.1098/rstb.2008.0006

Koenig, M. A., Clément, F., & Harris, P. L. (2004). Trust in testimony: Children's use of true and false statements. *Psychological Science*, *15*(10), 694–698.

Koenig, M. A., & Echols, C. H. (2003). Infants' understanding of false labeling events: The referential roles of words and the speakers who use them. *Cognition*, *87*(3), 179–208.

Koenig, M. A., & Harris, P. L. (2005). Preschoolers Mistrust Ignorant and Inaccurate Speakers. *Child Development*, *76*(6), 1261–1277. https://doi.org/10.1111/j.1467-8624.2005.00849.x

Koenig, M. A., & Harris, P. L. (2007). The basis of epistemic trust: Reliable testimony or reliable sources? *Episteme*, *4*(03), 264–284.

Kokal, I., Engel, A., Kirschner, S., & Keysers, C. (2011). Synchronized drumming enhances activity in the caudate and facilitates prosocial commitment - If the rhythm comes easily. *PLoS ONE*, *6*(11). Retrieved from http://www.scopus.com/inward/record.url?eid=2-s2.0-81155128181&partnerID=40&md5=524839ae0aad3ad77a93c6e55619d9bd

Koller, D., & Friedman, N. (2009). *Probabilistic Graphical Models: Principles and Techniques*. MIT Press.

Konvalinka, I., Xygalatas, D., Bulbulia, J., Schjodt, U., Jegindo, E.-M., Wallot, S., … Roepstorff, A. (2011). Synchronized arousal between performers and related spectators in a fire-walking ritual. *Proceedings of the National Academy of Sciences*, *108*(20), 8514–8519. https://doi.org/10.1073/pnas.1016955108

Konvalinka, Ivana, Bauer, M., Stahlhut, C., Hansen, L. K., Roepstorff, A., & Frith, C. D. (2014). Frontal alpha oscillations distinguish leaders from followers: Multivariate decoding of mutually interacting brains. *NeuroImage*, *94*, 79–88. https://doi.org/10.1016/j.neuroimage.2014.03.003

Konvalinka, Ivana, & Roepstorff, A. (2012). The two-brain approach: how can mutually interacting brains teach us something about social interaction? *Frontiers in Human Neuroscience*, *6*. https://doi.org/10.3389/fnhum.2012.00215

Kot, S., & Kulesza, W. (2016). The Chameleon as a Leech: The Costs of Mimicry for the Mimickee. *Polish Psychological Bulletin*, *47*(1), 131–135. https://doi.org/10.1515/ppb-2016-0014

Kouzakova, M., Karremans, J. C., van Baaren, R. B., & van Knippenberg, A. (2010). A Stranger's Cold Shoulder Makes the Heart Grow Fonder: Why Not Being Mimicked by a Stranger Enhances Longstanding Relationship Evaluations. *Social Psychological and Personality Science*, *1*(1), 87–93.

Kouzakova, M., van Baaren, R., & van Knippenberg, A. (2010). Lack of behavioral imitation in human interactions enhances salivary cortisol levels. *Hormones and Behavior*, *57*(4–5), 421–426.

Krumhuber, E., R, S., Cosker, D., Marshall, D., Rosin, P. L., & Kappas, A. (2007). Facial dynamics as indicators of trustworthiness and cooperative behavior. *Emotion*, *7*(4), 730–735. https://doi.org/10.1037/1528-3542.7.4.730

Kuhlen, A. K., & Brennan, S. E. (2013). Language in dialogue: when confederates might be hazardous to your data. *Psychonomic Bulletin & Review*, *20*(1), 54–72. https://doi.org/10.3758/s13423-012-0341-8

Kuhn, M. H., & McPartland, T. S. (1954). An empirical investigation of self-attitudes. *American Sociological Review*, *19*, 68–76. https://doi.org/10.2307/2088175

Kühn, S., Müller, B. C. N., van Baaren, R. B., Wietzker, A., Dijksterhuis, A., & Brass, M. (2010). Why do I like you when you behave like me? Neural mechanisms mediating positive consequences of observing someone being imitated. *Social Neuroscience*, *5*(4), 384–392. https://doi.org/10.1080/17470911003633750

Kulesza, W., Szypowska, Z., Jarman, M. S., & Dolinski, D. (2014). Attractive Chameleons Sell: The Mimicry-Attractiveness Link. *Psychology & Marketing*, *31*(7), 549–561. https://doi.org/10.1002/mar.20716

Kulms, P., Kopp, S., & Krämer, N. C. (2014). Let's Be Serious and Have a Laugh: Can Humor Support Cooperation with a Virtual Agent? In *Intelligent Virtual Agents* (pp. 250–259). Springer, Cham. https://doi.org/10.1007/978-3-319-09767-1_32

Kulms, P., Mattar, N., & Kopp, S. (2015). An Interaction Game Framework for the Investigation of Human–Agent Cooperation. In *Intelligent Virtual Agents* (pp. 399–402). Springer, Cham. https://doi.org/10.1007/978-3-319-21996-7_43

Kupper, Z., Ramseyer, F., Hoffmann, H., Kalbermatten, S., & Tschacher, W. (2010). Video-based quantification of body movement during social interaction indicates the severity of negative symptoms in patients with schizophrenia. *Schizophrenia Research*, *121*(1–3), 90–100. https://doi.org/10.1016/j.schres.2010.03.032

Kurzius, E., & Borkenau, P. (2015). Antecedents and Consequences of Mimicry: A Naturalistic Interaction Approach. *European Journal of Personality*, *29*(2), 107–124. https://doi.org/10.1002/per.1990

LaFrance, M. (1979). Nonverbal Synchrony and Rapport: Analysis by the Cross-Lag Panel Technique. *Social Psychology Quarterly*, *42*(1), 66. https://doi.org/10.2307/3033875

LaFrance, M. (1985). Postural Mirroring and Intergroup Relations. *Personality and Social Psychology Bulletin*, *11*(2), 207–217. https://doi.org/10.1177/0146167285112008

LaFrance, M., & Broadbent, M. (1976). Group Rapport: Posture Sharing as a Nonverbal Indicator. *Group & Organization Management*, *1*(3), 328–333. https://doi.org/10.1177/105960117600100307

Lakens, Daniel. (2013). Calculating and reporting effect sizes to facilitate cumulative science: a practical primer for t-tests and ANOVAs. *Cognition*, *4*, 863. https://doi.org/10.3389/fpsyg.2013.00863

Lakens, Daniël, Schubert, T., & Paladino, M.-P. (2016). Social Antecedents and Consequences of Behavioral Synchrony. In *Shared Representations: Sensorimotor Foundations of Social Life*.

Lakin, J. L., & Chartrand, T. L. (2003). Using Nonconscious Behavioral Mimicry to Create Affiliation and Rapport. *Psychological Science*, *14*(4), 334–339. https://doi.org/10.1111/1467-9280.14481

Lakin, J. L., Chartrand, T. L., & Arkin, R. M. (2008). I Am Too Just Like You: Nonconscious Mimicry as an Automatic Behavioral Response to Social Exclusion. *Psychological Science (Wiley-Blackwell)*, *19*(8), 816–822. https://doi.org/10.1111/j.1467-9280.2008.02162.x

Lakin, J. L., Jefferis, V. E., Cheng, C. M., & Chartrand, T. L. (2003). The Chameleon Effect as Social Glue: Evidence for the Evolutionary Significance of Nonconscious Mimicry. *Journal of Nonverbal Behavior*, *27*(3), 145–162. https://doi.org/10.1023/A:1025389814290

Lambdin, C., & Shaffer, V. A. (2009). Are within-subjects designs transparent? *Judgment and Decision Making*, *4*(7), 554–566.

Lamm, C., Batson, C., & Decety, J. (2007). The Neural Substrate of Human Empathy: Effects of Perspective-taking and Cognitive Appraisal. *Journal of Cognitive Neuroscience*, *19*(1), 42–58. https://doi.org/10.1162/jocn.2007.19.1.42

Landrum, A. R., Eaves Jr, B. S., & Shafto, P. (2015). Learning to trust and trusting to learn: a theoretical framework. *Trends in Cognitive Sciences*. https://doi.org/10.1016/j.tics.2014.12.007

Landrum, A. R., Mills, C. M., & Johnston, A. M. (2013). When do children trust the expert? Benevolence information influences children's trust more than expertise. *Developmental Science*, *16*(4), 622–638. https://doi.org/10.1111/desc.12059

Lane, J. D., Wellman, H. M., & Gelman, S. A. (2013). Informants' Traits Weigh Heavily in Young Children's Trust in Testimony and in Their Epistemic Inferences. *Child Development*, *84*(4), 1253–1268. https://doi.org/10.1111/cdev.12029

Launay, J., Dean, R. T., & Bailes, F. (2013). Synchronization Can Influence Trust Following Virtual Interaction. *Experimental Psychology (Formerly Zeitschrift Für Experimentelle Psychologie)*, *60*(1), 53–63. https://doi.org/10.1027/1618-3169/a000173

LaViola, J. J., Jr. (2000). A Discussion of Cybersickness in Virtual Environments. *SIGCHI Bull.*, *32*(1), 47–56. https://doi.org/10.1145/333329.333344

Leander, N. P., Chartrand, T. L., & Bargh, J. A. (2012). You Give Me the Chills Embodied Reactions to Inappropriate Amounts of Behavioral Mimicry. *Psychological Science*, *23*(7), 772–779. https://doi.org/10.1177/0956797611434535

Leander, N. P., Chartrand, T. L., & Wood, W. (2011). Mind your mannerisms: Behavioral mimicry elicits stereotype conformity. *Journal of Experimental Social Psychology*, *47*(1), 195–201. https://doi.org/10.1016/j.jesp.2010.09.002

Leighton, J., Bird, G., Orsini, C., & Heyes, C. (2010). Social attitudes modulate automatic imitation. *Journal of Experimental Social Psychology*, *46*(6), 905–910. https://doi.org/10.1016/j.jesp.2010.07.001

Li, B. Y. L., Mian, A. S., Liu, W., & Krishna, A. (2013). Using Kinect for face recognition under varying poses, expressions, illumination and disguise. In *2013 IEEE Workshop on Applications of Computer Vision (WACV)* (pp. 186–192). https://doi.org/10.1109/WACV.2013.6475017

Lis, S., Baer, N., Franzen, N., Hagenhoff, M., Gerlach, M., Koppe, G., … Kirsch, P. (2013). Social Interaction Behavior in ADHD in Adults in a Virtual Trust Game. *Journal of Attention Disorders*, 1087054713482581. https://doi.org/10.1177/1087054713482581

Liu, H., & Priest, S. (2009). Understanding public support for stem cell research: media communication, interpersonal communication and trust in key actors. *Public Understanding of Science*, *18*(6), 704–718. https://doi.org/10.1177/0963662508097625

Loomis, J. M., Blascovich, J. J., & Beall, A. C. (n.d.). Immersive virtual environment technology as a basic research tool in psychology. *Behavior Research Methods, Instruments, & Computers*, *31*(4), 557–564. https://doi.org/10.3758/BF03200735

Lumsden, J., Miles, L. K., & Macrae, C. N. (2012). Perceptions of Synchrony: Different Strokes for Different Folks? *Perception*, *41*(12), 1529–1531. https://doi.org/10.1068/p7360

Luschei, E. S., Ramig, L. O., Finnegan, E. M., Baker, K. K., & Smith, M. E. (2006). Patterns of Laryngeal Electromyography and the Activity of the Respiratory System During Spontaneous Laughter. *Journal of Neurophysiology*, *96*(1), 442–450. https://doi.org/10.1152/jn.00102.2006

Lyon, F., Möllering, G., & Saunders, M. (2012). *Handbook of Research Methods on Trust*. Edward Elgar Publishing.

MacDougall, H. G., & Moore, S. T. (2005). Marching to the beat of the same drummer: the spontaneous tempo of human locomotion. *Journal of Applied Physiology*, *99*(3), 1164–1173. https://doi.org/10.1152/japplphysiol.00138.2005

Macko, A., Malawski, M., & Tyszka, T. (2014). Belief in others' trustworthiness and trusting behaviour. *Polish Psychological Bulletin*, *45*(1), 43–51. https://doi.org/10.2478/ppb-2014-0007

Maddux, W. W., Mullen, E., & Galinsky, A. D. (2008). Chameleons bake bigger pies and take bigger pieces: Strategic behavioral mimicry facilitates negotiation outcomes. *Journal of Experimental Social Psychology*, *44*(2), 461–468. https://doi.org/10.1016/j.jesp.2007.02.003

Maister, L., Slater, M., Sanchez-Vives, M. V., & Tsakiris, M. (2015). Changing bodies changes minds: owning another body affects social cognition. *Trends in Cognitive Sciences*, *19*(1), 6–12. https://doi.org/10.1016/j.tics.2014.11.001

Manusov, V. (1992). Mimicry or synchrony: The effects of intentionality attributions for nonverbal mirroring behavior. *Communication Quarterly*, *40*(1), 69–83. https://doi.org/10.1080/01463379209369821

Markus, H. R., & Kitayama, S. (1991). Culture and the self: Implications for cognition, emotion, and motivation. *Psychological Review*, *98*(2), 224–253. https://doi.org/10.1037/0033-295X.98.2.224

Marsh, K. L., Richardson, M. J., & Schmidt, R. C. (2009). Social Connection Through Joint Action and Interpersonal Coordination. *Topics in Cognitive Science*, *1*(2), 320–339. https://doi.org/10.1111/j.1756-8765.2009.01022.x

Maurer, R. E., & Tindall, J. H. (1983). Effect of postural congruence on client's perception of counselor empathy. *Journal of Counseling Psychology*, *30*(2), 158–163. https://doi.org/10.1037/0022-0167.30.2.158

Maxwell, S. E., Lau, M. Y., & Howard, G. S. (2015). Is psychology suffering from a replication crisis? What does "failure to replicate" really mean? *American Psychologist*, *70*(6), 487–498. https://doi.org/10.1037/a0039400

McCall, C., & Blascovich, J. (2009a). How, When, and Why to Use Digital Experimental Virtual Environments to Study Social Behavior. *Social and Personality Psychology Compass*, *3*(5), 744–758. https://doi.org/10.1111/j.1751-9004.2009.00195.x

McCall, C., & Blascovich, J. (2009b). How, When, and Why to Use Digital Experimental Virtual Environments to Study Social Behavior. *Social and Personality Psychology Compass*, *3*(5), 744–758. https://doi.org/10.1111/j.1751-9004.2009.00195.x

McCall, C., Blascovich, J., Young, A., & Persky, S. (2009). Proxemic behaviors as predictors of aggression towards Black (but not White) males in an immersive virtual environment. *Social Influence*, *4*(2), 138–154. https://doi.org/10.1080/15534510802517418

McCall, C., Hildebrandt, L. K., Bornemann, B., & Singer, T. (2015). Physiophenomenology in retrospect: Memory reliably reflects physiological arousal during a prior threatening experience. *Consciousness and Cognition*, *38*, 60–70. https://doi.org/10.1016/j.concog.2015.09.011

McCall, C., Hildebrandt, L. K., Hartmann, R., Baczkowski, B. M., & Singer, T. (2016). Introducing the Wunderkammer as a tool for emotion research: Unconstrained gaze and movement patterns in three emotionally evocative virtual worlds.

Computers in Human Behavior, *59*, 93–107. https://doi.org/10.1016/j.chb.2016.01.028

McCall, C., & Singer, T. (2015). Facing Off with Unfair Others: Introducing Proxemic Imaging as an Implicit Measure of Approach and Avoidance during Social Interaction. *PLOS ONE*, *10*(2), e0117532. https://doi.org/10.1371/journal.pone.0117532

McCambridge, J., de Bruin, M., & Witton, J. (2012). The Effects of Demand Characteristics on Research Participant Behaviours in Non-Laboratory Settings: A Systematic Review. *PLoS ONE*, *7*(6), e39116. https://doi.org/10.1371/journal.pone.0039116

McClave, E. Z. (2000). Linguistic functions of head movements in the context of speech. *Journal of Pragmatics*, *32*(7), 855–878. https://doi.org/10.1016/S0378-2166(99)00079-X

McCroskey, J. C., & Teven, J. J. (1999). Goodwill: A reexamination of the construct and its measurement. *Communication Monographs*, *66*(1), 90–103. https://doi.org/10.1080/03637759909376464

McDowall, J. J. (1978). Interactional synchrony: A reappraisal. *Journal of Personality and Social Psychology*, *36*(9), 963–975. https://doi.org/10.1037/0022-3514.36.9.963

McEvily, B., Radzevick, J. R., & Weber, R. A. (2012). Whom do you distrust and how much does it cost? An experiment on the measurement of trust. *Games and Economic Behavior*, *74*(1), 285–298. https://doi.org/10.1016/j.geb.2011.06.011

McFarland, D. H. (2001). Respiratory Markers of Conversational Interaction. *Journal of Speech, Language, and Hearing Research*, *44*(1), 128–143. https://doi.org/10.1044/1092-4388(2001/012)

Mead, R., Atrash, A., & Matarić, M. J. (2013). Automated Proxemic Feature Extraction and Behavior Recognition: Applications in Human-Robot Interaction.

International Journal of Social Robotics, *5*(3), 367–378. https://doi.org/10.1007/s12369-013-0189-8

Meltzoff, A. N. (2007a). “Like me”: a foundation for social cognition. *Developmental Science*, *10*(1), 126–134. https://doi.org/10.1111/j.1467-7687.2007.00574.x

Meltzoff, A. N. (2007b). The “like me” framework for recognizing and becoming an intentional agent. *Acta Psychologica*, *124*(1), 26–43. https://doi.org/10.1016/j.actpsy.2006.09.005

Meltzoff, A. N., & Moore, M. K. (1977). Imitation of Facial and Manual Gestures by Human Neonates. *Science*, *198*(4312), 75–78.

Meltzoff, A. N., & Moore, M. K. (1983). Newborn Infants Imitate Adult Facial Gestures. *Child Development*, *54*(3), 702–709. https://doi.org/10.2307/1130058

Meltzoff, A. N., & Moore, M. K. (1997). Explaining Facial Imitation: A Theoretical Model. *Early Development & Parenting*, *6*(3–4), 179–192. https://doi.org/10.1002/(SICI)1099-0917(199709/12)6:3/4<179::AID-EDP157>3.0.CO;2-R

Messinger, D. S., Fogel, A., & Laurie, K. (1999). What’s in a smile? *Developmental Psychology*, *35*(3), 701–708. https://doi.org/10.1037/0012-1649.35.3.701

Messinger, D. S., Fogel, A., & Laurie, K. (2001). All smiles are positive, but some smiles are more positive than others. *Developmental Psychology*, *37*(5), 642–653. https://doi.org/10.1037/0012-1649.37.5.642

Michalski, R. S., Carbonell, J. G., & Mitchell, T. M. (2013). *Machine Learning: An Artificial Intelligence Approach*. Springer Science & Business Media.

Mihoub, A., Bailly, G., & Wolf, C. (2013). Social Behavior Modeling Based on Incremental Discrete Hidden Markov Models. In *Human Behavior Understanding* (pp. 172–183). Springer International Publishing. https://doi.org/10.1007/978-3-319-02714-2_15

Miles, L. K., Griffiths, J. L., Richardson, M. J., & Macrae, C. N. (2010). Too late to coordinate: Contextual influences on behavioral synchrony. *European Journal of Social Psychology*, *40*(1), 52–60. https://doi.org/10.1002/ejsp.721

Miles, L. K., Lumsden, J., Richardson, M. J., & Macrae, C. N. (2011). Do birds of a feather move together? Group membership and behavioral synchrony. *Experimental Brain Research*, *211*(3–4), 495–503. https://doi.org/10.1007/s00221-011-2641-z

Miles, L. K., Nind, L. K., & Macrae, C. N. (2009). The rhythm of rapport: Interpersonal synchrony and social perception. *Journal of Experimental Social Psychology*, *45*(3), 585–589. https://doi.org/10.1016/j.jesp.2009.02.002

Millar, W. S. (1988). Smiling, Vocal, and Attentive Behavior during Social Contingency Learning in Seven- and Ten-month-old Infants. *Merrill-Palmer Quarterly*, *34*(3), 301–325.

Miller, A. S., & Mitamura, T. (2003). Are Surveys on Trust Trustworthy? *Social Psychology Quarterly*, *66*(1), 62–70. https://doi.org/10.2307/3090141

Mills, C. M., Legare, C. H., Bills, M., & Mejias, C. (2010). Preschoolers use questions as a tool to acquire knowledge from different sources. *Journal of Cognition and Development*, *11*(4), 533–560.

Molden, D. C. (2014). Understanding Priming Effects in Social Psychology: What is “Social Priming” and How does it Occur? *Social Cognition*, *32*(Supplement), 1–11. https://doi.org/10.1521/soco.2014.32.supp.1

Molenberghs, P., Cunnington, R., & Mattingley, J. B. (2009). Is the mirror neuron system involved in imitation? A short review and meta-analysis. *Neuroscience & Biobehavioral Reviews*, *33*(7), 975–980. https://doi.org/10.1016/j.neubiorev.2009.03.010

Mondillon, L., Niedenthal, P. M., Gil, S., & Droit-Volet, S. (2007). Imitation of in-group versus out-group members’ facial expressions of anger: A test with a time perception task. *Social Neuroscience*, *2*(3–4), 223–237. https://doi.org/10.1080/17470910701376894

Morrill, R. J., Paukner, A., Ferrari, P. F., & Ghazanfar, A. A. (2012). Monkey lipsmacking develops like the human speech rhythm. *Developmental Science*, *15*(4), 557–568. https://doi.org/10.1111/j.1467-7687.2012.01149.x

Munhall, K. G., Jones, J. A., Callan, D. E., Kuratate, T., & Vatikiotis-Bateson, E. (2004). Visual Prosody and Speech Intelligibility Head Movement Improves Auditory Speech Perception. *Psychological Science*, *15*(2), 133–137. https://doi.org/10.1111/j.0963-7214.2004.01502010.x

Nadel, J. (2002). Imitation and imitation recognition: Functional use in preverbal infants and nonverbal children with autism. In *The Imitative Mind: Development, Evolution and Brain Bases*.

Nagaoka, C., & Komori, M. (2008). Body Movement Synchrony in Psychotherapeutic Counseling: A Study Using the Video-Based Quantification Method. *IEICE - Trans. Inf. Syst.*, *E91–D*(6), 1634–1640. https://doi.org/10.1093/ietisy/e91-d.6.1634

Nathoo, F. S., & Masson, M. E. J. (2016). Bayesian alternatives to null-hypothesis significance testing for repeated-measures designs. *Journal of Mathematical Psychology*, *72*, 144–157. https://doi.org/10.1016/j.jmp.2015.03.003

Navarre, D. D. (1982). Posture sharing in dyadic interaction. *American Journal of Dance Therapy*, *5*(1), 28–42. https://doi.org/10.1007/BF02579539

Neumann, R., & Strack, F. (2000). “Mood contagion”: the automatic transfer of mood between persons. *Journal of Personality and Social Psychology*, *79*(2), 211–223.

Nguyen, L. S., & Gatica-Perez, D. (2015). I Would Hire You in a Minute: Thin Slices of Nonverbal Behavior in Job Interviews. In *Proceedings of the 2015 ACM on International Conference on Multimodal Interaction* (pp. 51–58). New York, NY, USA: ACM. https://doi.org/10.1145/2818346.2820760

Niederhoffer, K. G., & Pennebaker, J. W. (2002). Linguistic Style Matching in Social Interaction. *Journal of Language and Social Psychology*, *21*(4), 337–360. https://doi.org/10.1177/026192702237953

Ninaus, M., Kober, S. E., Friedrich, E. V. C., Dunwell, I., Freitas, S. D., Arnab, S., … Neuper, C. (2014). Neurophysiological Methods for Monitoring Brain Activity in Serious Games and Virtual Environments: A Review. *Int. J. Technol. Enhanc. Learn.*, *6*(1), 78–103. https://doi.org/10.1504/IJTEL.2014.060022

Noah, J. A., Ono, Y., Nomoto, Y., Shimada, S., Tachibana, A., Zhang, X., … Hirsch, J. (2015). fMRI Validation of fNIRS Measurements During a Naturalistic Task. *Journal of Visualized Experiments : JoVE*, (100). https://doi.org/10.3791/52116

Noorden, L. van, & Moelants, D. (1999). Resonance in the Perception of Musical Pulse. *Journal of New Music Research*, *28*(1), 43–66. https://doi.org/10.1076/jnmr.28.1.43.3122

Oberman, L. M., & Ramachandran, V. S. (2007). The simulating social mind: The role of the mirror neuron system and simulation in the social and communicative deficits of autism spectrum disorders. *Psychological Bulletin*, *133*(2), 310–327. https://doi.org/10.1037/0033-2909.133.2.310

O'Doherty, J., Winston, J., Critchley, H., Perrett, D., Burt, D. M., & Dolan, R. J. (2003). Beauty in a smile: the role of medial orbitofrontal cortex in facial attractiveness. *Neuropsychologia*, *41*(2), 147–155.

Ohala, J. J. (1975). The temporal regulation of speech. In *Auditory Analysis and Perception of Speech*.

Oostenbroek, J., Suddendorf, T., Nielsen, M., Redshaw, J., Kennedy-Costantini, S., Davis, J., … Slaughter, V. (2016). Comprehensive Longitudinal Study Challenges the Existence of Neonatal Imitation in Humans. *Current Biology*, *26*(10), 1334–1338. https://doi.org/10.1016/j.cub.2016.03.047

Over, H., & Carpenter, M. (2009). Priming third-party ostracism increases affiliative imitation in children. *Developmental Science*, *12*(3), F1–F8. https://doi.org/10.1111/j.1467-7687.2008.00820.x

Pan, X., Banakou, D., & Slater, M. (2011). Computer Based Video and Virtual Environments in the Study of the Role of Emotions in Moral Behavior. In S. D'Mello, A. Graesser, B. Schuller, & J.-C. Martin (Eds.), *Affective Computing and Intelligent Interaction* (pp. 52–61). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-24571-8_6

Pan, X., Gillies, M., Barker, C., Clark, D. M., & Slater, M. (2012). Socially Anxious and Confident Men Interact with a Forward Virtual Woman: An Experimental Study. *PLoS ONE*, *7*(4), e32931. https://doi.org/10.1371/journal.pone.0032931

Pan, X., & Slater, M. (2011). Confronting a Moral Dilemma in Virtual Reality: A Pilot Study. In *Proceedings of the 25th BCS Conference on Human-Computer Interaction* (pp. 46–51). Swinton, UK, UK: British Computer Society. Retrieved from http://dl.acm.org/citation.cfm?id=2305316.2305326

Parrott, C. T., Brodsky, S. L., & Wilson, J. K. (2015). Thin slice expert testimony and mock trial deliberations. *International Journal of Law and Psychiatry*, *42–43*, 67–74. https://doi.org/10.1016/j.ijlp.2015.08.009

Pashler, H., Coburn, N., & Harris, C. R. (2012). Priming of Social Distance? Failure to Replicate Effects on Social and Food Judgments. *PLOS ONE*, *7*(8), e42510. https://doi.org/10.1371/journal.pone.0042510

Pasquini, E. S., Corriveau, K. H., Koenig, M., & Harris, P. L. (2007). Preschoolers monitor the relative accuracy of informants. *Developmental Psychology*, *43*(5), 1216–1226. https://doi.org/10.1037/0012-1649.43.5.1216

Paxton, A., & Dale, R. (2013). Frame-differencing methods for measuring bodily synchrony in conversation. *Behavior Research Methods*, *45*(2), 329–343. https://doi.org/10.3758/s13428-012-0249-2

Peck, T. C., Seinfeld, S., Aglioti, S. M., & Slater, M. (2013). Putting yourself in the skin of a black avatar reduces implicit racial bias. *Consciousness and Cognition*, *22*(3), 779–787. https://doi.org/10.1016/j.concog.2013.04.016

Pellegrini, R., & Ciceri, M. R. (2012). Listening to and mimicking respiration: Understanding and synchronizing joint actions. *Review of Psychology*, *19*(1), 17–27.

Pentland, A. (Sandy). (2010). *Honest Signals*. MIT Press.

Petty, R. E., Fazio, R. H., & Brinol, P. (2012). *Attitudes: Insights from the New Implicit Measures*. Psychology Press.

Piaget, J. (1946). *La formation du symbole chez enfant [Symbol formation in the child]*. Paris, France: Delachau & Nisette.

Pickering, M. J., & Garrod, S. (2013). An integrated theory of language production and comprehension. *Behavioral and Brain Sciences*, *36*(4), 329–347. https://doi.org/10.1017/S0140525X12001495

Piper, S. K., Krueger, A., Koch, S. P., Mehnert, J., Habermehl, C., Steinbrink, J., … Schmitz, C. H. (2014). A wearable multi-channel fNIRS system for brain imaging in freely moving subjects. *NeuroImage*, *85*, *Part 1*, 64–71. https://doi.org/10.1016/j.neuroimage.2013.06.062

Poppe, R., Zee, S., Heylen, D. K. J., & Taylor, P. J. (2013). AMAB: Automated measurement and analysis of body motion. *Behavior Research Methods*. https://doi.org/10.3758/s13428-013-0398-y

Ramseyer, F., & Tschacher, W. (2008). Synchrony in dyadic psychotherapy sessions. In *Simultaneity: Temporal structures and observer perspectives*.

Ramseyer, F., & Tschacher, W. (2011). Nonverbal synchrony in psychotherapy: Coordinated body movement reflects relationship quality and outcome. *Journal of Consulting and Clinical Psychology*, *79*(3), 284–295.

Ramseyer, Fabian, & Tschacher, W. (2010). Nonverbal Synchrony or Random Coincidence? How to Tell the Difference. In A. Esposito, N. Campbell, C. Vogel,

A. Hussain, & A. Nijholt (Eds.), *Development of Multimodal Interfaces: Active Listening and Synchrony* (pp. 182–196). Springer Berlin Heidelberg. Retrieved from http://link.springer.com/chapter/10.1007/978-3-642-12397-9_15

Ramseyer, Fabian, & Tschacher, W. (2014). Nonverbal synchrony of head- and body-movement in psychotherapy: different signals have different associations with outcome. *Psychology for Clinical Settings*, *5*, 979. https://doi.org/10.3389/fpsyg.2014.00979

Ravenscroft, I. (1998). What is it like to be someone else? Simulation and empathy. *Ratio*, *11*(2), 170–185.

Ray, E., & Heyes, C. (2011). Imitation in infancy: the wealth of the stimulus. *Developmental Science*, *14*(1), 92–105. https://doi.org/10.1111/j.1467-7687.2010.00961.x

Reddish, P., Fischer, R., & Bulbulia, J. (2013). Let’s Dance Together: Synchrony, Shared Intentionality and Cooperation. *PLoS ONE*, *8*(8). Retrieved from http://www.scopus.com/inward/record.url?eid=2-s2.0-84881343680&partnerID=40&md5=7624c9c4db847b738ccae7a5ee6fb8ee

Redeker, M., Stel, M., & Mastop, J. (2011). Does mimicking others change your self-view? *Journal of Social Psychology*, *151*(4), 387–390.

Reeves, B., & Nass, C. I. (1996). *The media equation: How people treat computers, television, and new media like real people and places* (Vol. xiv). New York, NY, US: Cambridge University Press.

Rempel, J. K., Holmes, J. G., & Zanna, M. P. (1985). Trust in close relationships. *Journal of Personality and Social Psychology*, *49*(1), 95–112. https://doi.org/10.1037/0022-3514.49.1.95

Rheingold, H. (1991). *Virtual Reality: Exploring the Brave New Technologies*. Simon & Schuster Adult Publishing Group.

Richardson, D. C., Dale, R., & Shockley, K. (2008). Synchrony and swing in conversation: coordination, temporal dynamics, and communication. In *Embodied Communication in Humans and Machines* (pp. 75–93). OUP Oxford.

Richardson, Daniel C., & Dale, R. (2005). Looking To Understand: The Coupling Between Speakers' and Listeners' Eye Movements and Its Relationship to Discourse Comprehension. *Cognitive Science*, *29*(6), 1045–1060. https://doi.org/10.1207/s15516709cog0000_29

Richardson, Daniel C., Dale, R., & Kirkham, N. Z. (2007). The Art of Conversation Is Coordination. *Psychological Science*, *18*(5), 407–413. https://doi.org/10.1111/j.1467-9280.2007.01914.x

Rizzolatti, G., & Craighero, L. (2004). The mirror-neuron system. *Annual Review of Neuroscience*, *27*, 169–192. https://doi.org/10.1146/annurev.neuro.27.070203.144230

Rosenberg, M. (1957). *Occupations and values*. Free Press.

Rotter, J. B. (1967). A new scale for the measurement of interpersonal trust1. *Journal of Personality*, *35*(4), 651–665. https://doi.org/10.1111/j.1467-6494.1967.tb01454.x

Rotter, J. B. (1971). Generalized expectancies for interpersonal trust. *American Psychologist*, *26*(5), 443–452. https://doi.org/10.1037/h0031464

Rousseau, D. M., Sitkin, S. B., Burt, R. S., & Camerer, C. (1998). Not So Different After All: A Cross-Discipline View Of Trust. *Academy of Management Review*, *23*(3), 393–404. https://doi.org/10.5465/AMR.1998.926617

Ruby, P., & Decety, J. (2001). Effect of subjective perspective taking during simulation of action: a PET investigation of agency. *Nature Neuroscience*, *4*(5), 546–550. https://doi.org/10.1038/87510

Sanchez-Burks, J., Bartel, C. A., & Blount, S. (2009). Performance in intercultural interactions at work: Cross-cultural differences in response to behavioral mirroring. *Journal of Applied Psychology*, *94*(1), 216–223. https://doi.org/10.1037/a0012829

Sanchez-Vives, M. V., & Slater, M. (2005). From presence to consciousness through virtual reality. *Nature Reviews Neuroscience*, *6*(4), 332–339. https://doi.org/10.1038/nrn1651

Sanfey, A. G., Rilling, J. K., Aronson, J. A., Nystrom, L. E., & Cohen, J. D. (2003). The neural basis of economic decision-making in the Ultimatum Game. *Science (New York, N.Y.)*, *300*(5626), 1755–1758. https://doi.org/10.1126/science.1082976

Saxe, R. (2009). The neural evidence for simulation is weaker than i think you think it is. *Philosophical Studies*, *144*(3), 447–456.

Schechter, L. (2007). Traditional trust measurement and the risk confound: An experiment in rural Paraguay. *Journal of Economic Behavior & Organization*, *62*(2), 272–292. https://doi.org/10.1016/j.jebo.2005.03.006

Scheflen, A. E. (1964). The significance of posture in communication systems. *Psychiatry: Journal for the Study of Interpersonal Processes*, *27*(4), 316–331.

Scheflen, A. E. (1972). Body Language and the Social Order; Communication as Behavioral Control. Retrieved from http://eric.ed.gov/?id=ED077060

Schilbach, L., Eickhoff, S. B., Cieslik, E., Shah, N. J., Fink, G. R., & Vogeley, K. (2011). Eyes on me: an fMRI study of the effects of social gaze on action control. *Social Cognitive and Affective Neuroscience*, *6*(4), 393–403. https://doi.org/10.1093/scan/nsq067

Schilbach, L., Wilms, M., Eickhoff, S. B., Romanzetti, S., Tepest, R., Bente, G., … Vogeley, K. (2009). Minds Made for Sharing: Initiating Joint Attention Recruits Reward-related Neurocircuitry. *Journal of Cognitive Neuroscience*, *22*(12), 2702–2715. https://doi.org/10.1162/jocn.2009.21401

Schilbach, L., Wilms, M., Eickhoff, S. B., Romanzetti, S., Tepest, R., Bente, G., … Vogeley, K. (2010). Minds made for sharing: initiating joint attention recruits reward-related neurocircuitry. *Journal of Cognitive Neuroscience*, *22*(12), 2702–2715. https://doi.org/10.1162/jocn.2009.21401

Schmidt, R. C., Morr, S., Fitzpatrick, P., & Richardson, M. J. (2012). Measuring the Dynamics of Interactional Synchrony. *Journal of Nonverbal Behavior*, *36*(4), 263–279. https://doi.org/10.1007/s10919-012-0138-5

Schmidt, R. C., Nie, L., Franco, A., & Richardson, M. J. (2014). Bodily synchronization underlying joke telling. *Frontiers in Human Neuroscience*, *8*, 633. https://doi.org/10.3389/fnhum.2014.00633

Schniter, E., Sheremeta, R. M., & Shields, T. W. (2013). *Limitations to Signaling Trust with All or Nothing Investments* (Working Paper No. 13–24). Chapman University, Economic Science Institute. Retrieved from https://ideas.repec.org/p/chu/wpaper/13-24.html

Shanks, D. R., Newell, B. R., Lee, E. H., Balakrishnan, D., Ekelund, L., Cenac, Z., … Moore, C. (2013). Priming Intelligent Behavior: An Elusive Phenomenon. *PLOS ONE*, *8*(4), e56515. https://doi.org/10.1371/journal.pone.0056515

Shockley, K., Baker, A. A., Richardson, M. J., & Fowler, C. A. (2007). Articulatory constraints on interpersonal postural coordination. *Journal of Experimental Psychology: Human Perception and Performance*, *33*(1), 201–208. https://doi.org/10.1037/0096-1523.33.1.201

Shockley, K., Butwill, M., Zbilut, J. P., & Webber, C. L. (2002). Cross recurrence quantification of coupled oscillators. *Physics Letters A*, *305*(1), 59–69. https://doi.org/10.1016/S0375-9601(02)01411-1

Shockley, K., Richardson, D. C., & Dale, R. (2009). Conversation and Coordinative Structures. *Topics in Cognitive Science*, *1*(2), 305–319. https://doi.org/10.1111/j.1756-8765.2009.01021.x

Shockley, K., Santana, M.-V., & Fowler, C. A. (2003). Mutual interpersonal postural constraints are involved in cooperative conversation. *Journal of Experimental*

Psychology: Human Perception and Performance, *29*(2), 326–332. https://doi.org/10.1037/0096-1523.29.2.326

Slater, M., Pertaub, D. P., & Steed, A. (1999). Public speaking in virtual reality: facing an audience of avatars. *IEEE Computer Graphics and Applications*, *19*(2), 6–9. https://doi.org/10.1109/38.749116

Slater, Mel, Antley, A., Davison, A., Swapp, D., Guger, C., Barker, C., … Sanchez-Vives, M. V. (2006). A Virtual Reprise of the Stanley Milgram Obedience Experiments. *PLOS ONE*, *1*(1), e39. https://doi.org/10.1371/journal.pone.0000039

Slater, Mel, Pérez Marcos, D., Ehrsson, H., Sanchez-Vives, M. V., Slater, M., Perez-Marcos, D., … Sanchez-Vives, M. V. (2009). Inducing illusory ownership of a virtual body. *Frontiers in Neuroscience*, *3*. https://doi.org/10.3389/neuro.01.029.2009

Slater, Mel, Pertaub, D.-P., Barker, C., & Clark, D. M. (2006). An Experimental Study on Fear of Public Speaking Using a Virtual Environment. *CyberPsychology & Behavior*, *9*(5), 627–633. https://doi.org/10.1089/cpb.2006.9.627

Slater, Mel, Rovira, A., Southern, R., Swapp, D., Zhang, J. J., Campbell, C., & Levine, M. (2013). Bystander Responses to a Violent Incident in an Immersive Virtual Environment. *PLOS ONE*, *8*(1), e52766. https://doi.org/10.1371/journal.pone.0052766

Smith, A. (1759). *The theory of moral sentiments*. London, UK: Penguin Books Ltd.

Sofianidis, G., Elliott, M. T., Wing, A. M., & Hatzitaki, V. (2014). Can dancers suppress the haptically mediated interpersonal entrainment during rhythmic sway? *Acta Psychologica*, *150*, 106–113. https://doi.org/10.1016/j.actpsy.2014.05.002

Sorrentino, R. M., Holmes, J. G., Hanna, S. E., & Sharp, A. (1995). Uncertainty orientation and trust in close relationships: Individual differences in cognitive styles. *Journal of Personality and Social Psychology*, *68*(2), 314–327. https://doi.org/10.1037/0022-3514.68.2.314

Sparenberg, P., Topolinski, S., Springer, A., & Prinz, W. (2012). Minimal mimicry: Mere effector matching induces preference. *Brain and Cognition*, *80*(3), 291–300. https://doi.org/10.1016/j.bandc.2012.08.004

Spaulding, S. (2012). Mirror neurons are not evidence for the Simulation Theory. *Synthese*, *189*(3), 515–534. https://doi.org/10.1007/s11229-012-0086-y

Spengler, S., von Cramon, D. Y., & Brass, M. (2010). Resisting motor mimicry: Control of imitation involves processes central to social cognition in patients with frontal and temporo-parietal lesions. *Social Neuroscience*, *5*(4), 401–416.

Spengler, Stephanie, von Cramon, D. Y., & Brass, M. (2009). Control of shared representations relies on key processes involved in mental state attribution. *Human Brain Mapping*, *30*(11), 3704–3718. https://doi.org/10.1002/hbm.20800

Stanley, D. A., Sokol-Hessner, P., Fareri, D. S., Perino, M. T., Delgado, M. R., Banaji, M. R., & Phelps, E. A. (2012). Race and reputation: perceived racial group trustworthiness influences the neural correlates of trust decisions. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, *367*(1589), 744–753. https://doi.org/10.1098/rstb.2011.0300

Stel, M., Blascovich, J., McCall, C., Mastop, J., van Baaren, R. B., & Vonk, R. (2010). Mimicking disliked others: Effects of a priori liking on the mimicry-liking link. *European Journal of Social Psychology*, *40*(5), 867–880. https://doi.org/10.1002/ejsp.655

Stel, M., Dijk, E. van, & Olivier, E. (2009). You Want to Know the Truth? Then Don’t Mimic! *Psychological Science*, *20*(6), 693–699. https://doi.org/10.1111/j.1467-9280.2009.02350.x

Stel, M., & Harinck, F. (2011). Being Mimicked Makes You a Prosocial Voter. *Experimental Psychology (Formerly Zeitschrift Für Experimentelle Psychologie)*, *58*(1), 79–84. https://doi.org/10.1027/1618-3169/a000070

Stel, M., Rispens, S., Leliveld, M., & Lokhorst, A. M. (2011). The consequences of mimicry for prosocials and proselfs: Effects of social value orientation on the mimicry-liking link. *European Journal of Social Psychology*, *41*(3), 269–274. https://doi.org/10.1002/ejsp.790

Stel, M., van Baaren, R. B., Blascovich, J., van Dijk, E., McCall, C., Pollmann, M. M. H., … Vonk, R. (2010). Effects of a Priori Liking on the Elicitation of Mimicry. *Experimental Psychology (Formerly Zeitschrift Für Experimentelle Psychologie)*, *57*(6), 412–418. https://doi.org/10.1027/1618-3169/a000050

Stel, M., van den Heuvel, C., & Smeets, R. C. (2008). Facial feedback mechanisms in autistic spectrum disorders. *Journal of Autism and Developmental Disorders*, *38*(7), 1250–1258. https://doi.org/10.1007/s10803-007-0505-y

Stel, M., & Vonk, R. (2010). Mimicry in social interaction: Benefits for mimickers, mimickees, and their interaction. *British Journal of Psychology*, *101*(2), 311–323. https://doi.org/10.1348/000712609X465424

Stel, M., van Baaren, R. B., Blascovich, J., van Dijk, E., McCall, C., Pollmann, M. M. H., … Vonk, R. (2010). Effects of a Priori Liking on the Elicitation of Mimicry. *Experimental Psychology*, *57*(6), 412–418. https://doi.org/10.1027/1618-3169/a000050

Stel, M., van Dijk, E., & van Baaren, R. B. (2016). When and Why Mimicry is Facilitated and Attenuated. *Social and Personality Psychology Compass*, *10*(10), 561–574. https://doi.org/10.1111/spc3.12269

Stich, S., & Nichols, S. (1992). Folk Psychology: Simulation or Tacit Theory? *Mind & Language*, *7*(1–2), 35–71. https://doi.org/10.1111/j.1468-0017.1992.tb00196.x

Stroebe, W., & Strack, F. (2014). The Alleged Crisis and the Illusion of Exact Replication. *Perspectives on Psychological Science*, *9*(1), 59–71. https://doi.org/10.1177/1745691613514450

Szuster, A., & Wojnarowska, A. (2016). The Influence of Mimicry on the Reduction of Infra-Humanization. *Frontiers in Psychology*, *7*. https://doi.org/10.3389/fpsyg.2016.00975

Tanner, R. (2008). The Convincing Chameleon: Impact of Mimicry on Persuasion. *The Journal of Consumer Research*, *34*, 754.

Tanner, R. J., Ferraro, R., Chartrand, T. L., Bettman, J. R., & Baaren, R. V. (2008). Of Chameleons and Consumption: The Impact of Mimicry on Choice and Preferences. *Journal of Consumer Research*, *34*(6), 754–766. https://doi.org/10.1086/522097

Thorndike, E. (1898). *Animal Intelligence: Experimental Studies*. London, UK: Macmillan & Co., Ltd.

Todorov, A. (2008). Evaluating Faces on Trustworthiness. *Annals of the New York Academy of Sciences*, *1124*(1), 208–224. https://doi.org/10.1196/annals.1440.012

Todorov, A., & Duchaine, B. (2008). Reading trustworthiness in faces without recognizing faces. *Cognitive Neuropsychology*, *25*(3), 395–410. https://doi.org/10.1080/02643290802044996

Todorov, A., Pakrashi, M., & Oosterhof, N. N. (2009). Evaluating Faces on Trustworthiness After Minimal Time Exposure. *Social Cognition*, *27*(6), 813–833. https://doi.org/10.1521/soco.2009.27.6.813

Tortosa, María I., Lupiáñez, J., & Ruz, M. (2013). Race, emotion and trust: An ERP study. *Brain Research*, *1494*, 44–55. https://doi.org/10.1016/j.brainres.2012.11.037

Tortosa, Maria I., Strizhko, T., Capizzi, M., & Ruz, M. (2013). Interpersonal Effects of Emotion in a Multi-Round Trust Game. *Psicologica: International Journal of Methodology and Experimental Psychology*, *34*(2), 179–198.

Tressoldi, P. E. (2012). Replication Unreliability in Psychology: Elusive Phenomena or “Elusive” Statistical Power? *Frontiers in Psychology*, *3*. https://doi.org/10.3389/fpsyg.2012.00218

Tsai, J. C.-C., Sebanz, N., & Knoblich, G. (2011). The GROOP effect: Groups mimic group actions. *Cognition*, *118*(1), 135–140. https://doi.org/10.1016/j.cognition.2010.10.007

Tzieropoulos, Hélène. (2013). The Trust Game in neuroscience: A short review. *Social Neuroscience*, *8*(5), 407–416. https://doi.org/10.1080/17470919.2013.832375

Tzieropoulos, Helène, Grave De Peralta, R., Bossaerts, P., & Gonzalez Andino, S. L. (2011). The impact of disappointment in decision making: inter-individual differences and electrical neuroimaging. *Frontiers in Human Neuroscience*, *4*, 235. https://doi.org/10.3389/fnhum.2010.00235

Uddin, L. Q. (2015). Salience processing and insular cortical function and dysfunction. *Nature Reviews Neuroscience*, *16*(1), 55–61. https://doi.org/10.1038/nrn3857

Uddin, L. Q., Molnar-Szakacs, I., Zaidel, E., & Iacoboni, M. (2006). rTMS to the right inferior parietal lobule disrupts self–other discrimination. *Social Cognitive and Affective Neuroscience*, *1*(1), 65–71. https://doi.org/10.1093/scan/nsl003

Valdesolo, P., & DeSteno, D. (2011). Synchrony and the social tuning of compassion. *Emotion*, *11*(2), 262–266.

Valdesolo, Piercarlo, Ouyang, J., & DeSteno, D. (2010). The rhythm of joint action: Synchrony promotes cooperative ability. *Journal of Experimental Social Psychology*, *46*(4), 693–695. https://doi.org/10.1016/j.jesp.2010.03.004

Van Baaren, R. B., & Chartrand, T. L. (2005). Nonconscious imitation has consequences that go beyond the dyad (pp. 128–132). Retrieved from http://www.scopus.com/inward/record.url?eid=2-s2.0-84858838776&partnerID=40&md5=594cff6949e9b3c3a9f00be123b0148e

van Baaren, R. B., Decety, J., Dijksterhuis, A., van der Leij, A., & van Leeuwen, M. L. (2009). Being imitated: Consequences of nonconsciously showing empathy. In J. Decety & W. Ickes (Eds.), *The social neuroscience of empathy* (pp. 31–42). Cambridge, MA, US: MIT Press.

van Baaren, R. B., Holland, R. W., Kawakami, K., & van Knippenberg, A. (2004). Mimicry and Prosocial Behavior. *Psychological Science*, *15*(1), 71–74. https://doi.org/10.1111/j.0963-7214.2004.01501012.x

van Baaren, R. B., Janssen, L., Chartrand, T. L., & Dijksterhuis, A. (2009). Where is the love? The social aspects of mimicry. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *364*(1528), 2381–2389. https://doi.org/10.1098/rstb.2009.0057

van Baaren, R. B., Maddux, W. W., Chartrand, T. L., de Bouter, C., & van Knippenberg, A. (2003). It takes two to mimic: Behavioral consequences of self-construals. *Journal of Personality and Social Psychology*, *84*(5), 1093–1102. https://doi.org/10.1037/0022-3514.84.5.1093

van der Schalk, J., Fischer, A., Doosje, B., Wigboldus, D., Hawk, S., Rotteveel, M., & Hess, U. (2011). Convergent and divergent responses to emotional displays of ingroup and outgroup. *Emotion*, *11*(2), 286–298. https://doi.org/10.1037/a0022582

van Swol, L. M. (2003). The Effects of Nonverbal Mirroring on Perceived Persuasiveness, Agreement with an Imitator, and Reciprocity in a Group Discussion. *Communication Research*, *30*(4), 461–480. https://doi.org/10.1177/0093650203253318

van 't Veer, A. E., & Giner-Sorolla, R. (2016). Pre-registration in social psychology—A discussion and suggested template. *Journal of Experimental Social Psychology*, *67*, 2–12. https://doi.org/10.1016/j.jesp.2016.03.004

VanderBorght, M., & Jaswal, V. K. (2009). Who Knows Best? Preschoolers Sometimes Prefer Child Informants Over Adult Informants. *Infant and Child Development*, *18*(1), 61–71. https://doi.org/10.1002/icd.591

Varlet, M., Marin, L., Lagarde, J., & Bardy, B. G. (2011). Social postural coordination. *Journal of Experimental Psychology: Human Perception and Performance*, *37*(2), 473–483. https://doi.org/10.1037/a0020552

Verberne, F. M. F., Ham, J., & Midden, C. J. H. (2015). Trusting a Virtual Driver That Looks, Acts, and Thinks Like You. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 0018720815580749. https://doi.org/10.1177/0018720815580749

Verberne, F. M. F., Ham, J., Ponnada, A., & Midden, C. J. H. (2013a). Trusting Digital Chameleons: The Effect of Mimicry by a Virtual Social Agent on User Trust. In S. Berkovsky & J. Freyne (Eds.), *Persuasive Technology* (pp. 234–245). Springer Berlin Heidelberg. Retrieved from http://link.springer.com/chapter/10.1007/978-3-642-37157-8_28

Verberne, F. M. F., Ham, J., Ponnada, A., & Midden, C. J. H. (2013b). *Trusting digital chameleons: The effect of mimicry by a virtual social agent on user trust* (Vol. 7822 LNCS).

Vohs, K. D., Mead, N. L., & Goode, M. R. (2006). The Psychological Consequences of Money. *Science*, *314*(5802), 1154–1156. https://doi.org/10.1126/science.1132491

Vrijsen, J. N., Lange, W.-G., Becker, E. S., & Rinck, M. (2010). Socially anxious individuals lack unintentional mimicry. *Behaviour Research and Therapy*, *48*(6), 561–564. https://doi.org/10.1016/j.brat.2010.02.004

Vrijsen, J. N., Lange, W.-G., Dotsch, R., Wigboldus, D. H. J., & Rinck, M. (2010). How do socially anxious women evaluate mimicry? A virtual reality study. *Cognition & Emotion*, *24*(5), 840–847. https://doi.org/10.1080/13854040902833652

Wagenmakers, E.-J. (2007). A practical solution to the pervasive problems ofp values. *Psychonomic Bulletin & Review*, *14*(5), 779–804. https://doi.org/10.3758/BF03194105

Wallbott, H. G. (1995). Congruence, contagion, and motor mimicry: Mutualities in nonverbal exchange. In I. Markova, C. F. Graumann, & K. Foppa (Eds.), *Mutualities in dialogue* (pp. 82–98). Cambridge University Press.

Walton, A. E., Richardson, M. J., Langland-Hassan, P., & Chemero, A. (2015). Improvisation and the self-organization of multiple musical bodies. *Frontiers in Psychology*, *6*. https://doi.org/10.3389/fpsyg.2015.00313

Wang, Y., & Hamilton, A. F. de C. (2012). Social top-down response modulation (STORM): a model of the control of mimicry in social interaction. *Frontiers in Human Neuroscience*, *6*. https://doi.org/10.3389/fnhum.2012.00153

Wang, Y., & Hamilton, A. F. de C. (2013). Understanding the Role of the "Self" in the Social Priming of Mimicry. *PLoS ONE*, *8*(4), e60249. https://doi.org/10.1371/journal.pone.0060249

Wang, Y., Newport, R., & Hamilton, A. F. de C. (2011). Eye contact enhances mimicry of intransitive hand movements. *Biology Letters*, *7*(1), 7–10. https://doi.org/10.1098/rsbl.2010.0279

Wang, Y., Ramsey, R., & Hamilton, A. F. de C. (2011). The Control of Mimicry by Eye Contact Is Mediated by Medial Prefrontal Cortex. *The Journal of Neuroscience*, *31*(33), 12001–12010. https://doi.org/10.1523/JNEUROSCI.0845-11.2011

Warner, R. M. (1996). Coordinated cycles in behavior and physiology during face-to-face social interactions. In J. H. Watt & C. A. VanLear (Eds.), *Dynamic patterns in communication processes* (pp. 327–352). Thousand Oaks, CA, US: Sage Publications, Inc.

Washburn, A., DeMarco, M., de Vries, S., Ariyabuddhiphongs, K., Schmidt, R. C., Richardson, M. J., & Riley, M. A. (2014). Dancers entrain more effectively than non-dancers to another actor's movements. *Frontiers in Human Neuroscience*, *8*, 800. https://doi.org/10.3389/fnhum.2014.00800

Watson, J. S., & Ramey, C. T. (1972). Reactions to response-contingent stimulation in infancy. *Merrill-Palmer Quarterly of Behavior and Development*, *18*(3), 219–227.

Whiten, A., Horner, V., & Marshall-Pescini, S. (2005). Selective Imitation in Child and Chimpanzee: A Window on the Construal of Others' Actions. In *Perspectives on Imitation: Mechanisms of imitation and imitation in animals*.

Whiten, Andrew, McGuigan, N., Marshall-Pescini, S., & Hopper, L. M. (2009). Emulation, imitation, over-imitation and the scope of culture for child and chimpanzee. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, *364*(1528), 2417–2428. https://doi.org/10.1098/rstb.2009.0069

Williams, L. E., & Bargh, J. A. (2008). Keeping One's Distance: The Influence of Spatial Distance Cues on Affect and Evaluation. *Psychological Science*, *19*(3), 302–308. https://doi.org/10.1111/j.1467-9280.2008.02084.x

Willinger, M., Keser, C., Lohmann, C., & Usunier, J.-C. (2003). A comparison of trust and reciprocity between France and Germany: Experimental investigation based on the investment game. *Journal of Economic Psychology*, *24*(4), 447–466. https://doi.org/10.1016/S0167-4870(02)00165-4

Willis, J., & Todorov, A. (2006). First Impressions Making Up Your Mind After a 100-Ms Exposure to a Face. *Psychological Science*, *17*(7), 592–598. https://doi.org/10.1111/j.1467-9280.2006.01750.x

Wilms, M., Schilbach, L., Pfeiffer, U., Bente, G., Fink, G. R., & Vogeley, K. (2010). It's in your eyes—using gaze-contingent stimuli to create truly interactive paradigms for. *Social Cognitive and Affective Neuroscience*, *5*(1), 98–107. https://doi.org/10.1093/scan/nsq024

Wiltermuth, S. S., & Heath, C. (2009). Synchrony and cooperation. *Psychological Science*, *20*(1), 1–5.

Won, A. S., Bailenson, J. N., Stathatos, S. C., & Dai, W. (2014). Automatically Detected Nonverbal Behavior Predicts Creativity in Collaborating Dyads. *Journal of Nonverbal Behavior*, *38*(3), 389–408. https://doi.org/10.1007/s10919-014-0186-0

Yabar, Y., Johnston, L., Miles, L., & Peace, V. (2006). Implicit behavioral mimicry: Investigating the impact of group membership. *Journal of Nonverbal Behavior*, *30*(3), 97–113.

Yamaoka, F., Kanda, T., Ishiguro, H., & Hagita, N. (2006). How Contingent Should a Communication Robot Be? In *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-robot Interaction* (pp. 313–320). New York, NY, USA: ACM. https://doi.org/10.1145/1121241.1121294

Yamaoka, F., Kanda, T., Ishiguro, H., & Hagita, N. (2007). How contingent should a lifelike robot be? The relationship between contingency and complexity. *Connection Science*, *19*(2), 143–162. https://doi.org/10.1080/09540090701371519

Yong, E. (2012). Nobel laureate challenges psychologists to clean up their act. *Nature News*. https://doi.org/10.1038/nature.2012.11535

Zaitsev, M., Dold, C., Sakas, G., Hennig, J., & Speck, O. (2006). Magnetic resonance imaging of freely moving objects: prospective real-time motion correction using an external optical motion tracking system. *NeuroImage*, *31*(3), 1038–1050. https://doi.org/10.1016/j.neuroimage.2006.01.039

Zanbaka, C., Goolkasian, P., & Hodges, L. (2006). Can a Virtual Cat Persuade You?: The Role of Gender and Realism in Speaker Persuasiveness. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 1153–1162). New York, NY, USA: ACM. https://doi.org/10.1145/1124772.1124945

Zbilut, J. P., Giuliani, A., & Webber, C. L. (1998a). Detecting deterministic signals in exceptionally noisy environments using cross-recurrence quantification. *Physics Letters A*, *246*(1), 122–128. https://doi.org/10.1016/S0375-9601(98)00457-5

Zbilut, J. P., Giuliani, A., & Webber, C. L. (1998b). Recurrence quantification analysis and principal components in the detection of short complex signals. *Physics Letters A*, *237*(3), 131–135. https://doi.org/10.1016/S0375-9601(97)00843-8